



US011715464B2

(12) **United States Patent**
Nickson et al.

(10) **Patent No.:** **US 11,715,464 B2**

(45) **Date of Patent:** **Aug. 1, 2023**

(54) **USING AUGMENTATION TO CREATE NATURAL LANGUAGE MODELS**

G06F 40/211; G06F 40/216; G06F 40/247; G06F 40/284; G06F 40/35; G06F 16/90332; G06F 40/56

(71) Applicant: **Apple Inc.**, Cupertino, CA (US)

See application file for complete search history.

(72) Inventors: **Thomas Robert Nickson**, Campbell, CA (US); **Keith Scott Brisson**, San Francisco, CA (US); **Eric Gregory**, Larkspur, CA (US); **Thomas B. Gunter**, Oxford (GB); **Arthur A. Van Hoff**, Menlo Park, CA (US)

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,156,868 A * 5/1979 Levinson G10L 15/193 704/256.4

9,380,155 B1 6/2016 Reding et al.
2009/0048841 A1 2/2009 Pollet et al.
2011/0238407 A1 9/2011 Kent

(Continued)

(73) Assignee: **Apple Inc.**, Cupertino, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

OTHER PUBLICATIONS

Nakagawa, T., Kudo, T., & Matsumoto, Y. (2001). Unknown word guessing and part-of-speech tagging using support vector machines. In NLP RS (pp. 325-331). (Year: 2001).*

(Continued)

(21) Appl. No.: **17/366,519**

(22) Filed: **Jul. 2, 2021**

Primary Examiner — Bhavesh M Mehta

Assistant Examiner — Philip H Lam

(65) **Prior Publication Data**

US 2022/0084511 A1 Mar. 17, 2022

(74) *Attorney, Agent, or Firm* — DLA Piper LLP (US)

(57) **ABSTRACT**

Systems and processes for creating and updating natural language models are provided. An example process of creating a natural language model includes, at an electronic device with one or more processors and memory, receiving an utterance, associating an action structure with the utterance, determining a plurality of augmented utterances based on the received utterance, creating a natural language model including the received utterance and the plurality of augmented utterances by mapping the plurality of augmented utterance to the associated action structure, and providing the natural language model including the received utterance and the plurality of augmented utterances to a second electronic device.

42 Claims, 16 Drawing Sheets

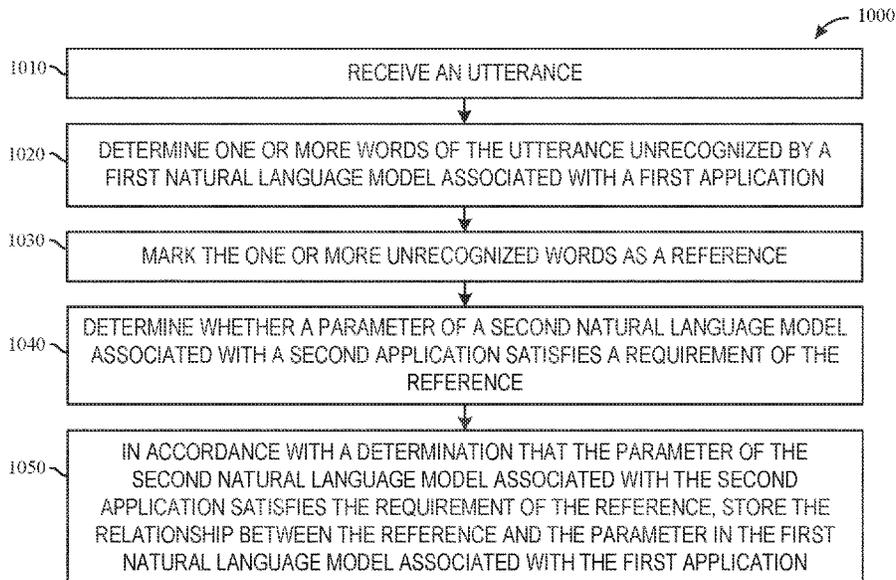
Related U.S. Application Data

(60) Provisional application No. 63/078,280, filed on Sep. 14, 2020.

(51) **Int. Cl.**
G10L 15/183 (2013.01)
G10L 15/06 (2013.01)

(52) **U.S. Cl.**
CPC **G10L 15/183** (2013.01); **G10L 15/063** (2013.01); **G10L 2015/0635** (2013.01)

(58) **Field of Classification Search**
CPC G10L 15/183; G10L 15/063; G10L 2015/0635; G10L 15/1822; G06F 40/157;



(56)

References Cited

U.S. PATENT DOCUMENTS

2012/0016678	A1*	1/2012	Gruber	G06N 5/041 704/E21.001
2014/0112556	A1	4/2014	Kalinli-Akbacak	
2014/0198048	A1	7/2014	Unruh et al.	
2019/0237061	A1*	8/2019	Rusak	G10L 15/063
2019/0272818	A1	9/2019	Fernandez et al.	
2019/0361978	A1*	11/2019	Ray	G06F 40/30
2021/0082400	A1*	3/2021	Vishnoi	G06F 40/35
2021/0304075	A1*	9/2021	Duong	G06F 40/35

OTHER PUBLICATIONS

Alsharif et al., "Long Short-Term Memory Neural Network for Keyboard Gesture Decoding", IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Brisbane, Australia, Sep. 2015, 5 pages.

Berry et al., "PTIME: Personalized Assistance for Calendaring", ACM Transactions on Intelligent Systems and Technology, vol. 2, No. 4, Article 40, Jul. 2011, pp. 1-22.

Büttner et al., "The Design Space of Augmented and Virtual Reality Applications for Assistive Environments in Manufacturing: A Visual Approach", In Proceedings of the 10th International Conference on Pervasive Technologies Related to Assistive Environments (PETRA '17), Island of Rhodes, Greece, Online available at: <https://dl.acm.org/doi/pdf/10.1145/3056540.3076193>, Jun. 21-23, 2017, pp. 433-440.

Chen et al., "Progressive Joint Modeling in Unsupervised Single-Channel Overlapped Speech Recognition", IEEE/ACM Transactions on Audio, Speech, and Language Processing, vol. 26, No. 1, Jan. 2018, pp. 184-196.

Conneau et al., "Supervised Learning of Universal Sentence Representations from Natural Language Inference Data", Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, Copenhagen, Denmark, Sep. 7-11, 2017, pp. 670-680.

Coulouris et al., "Distributed Systems: Concepts and Design (Fifth Edition)", Addison-Wesley, May 7, 2011, 391 pages.

Dighe et al., "Lattice-based improvements for voice triggering using graph neural networks", in 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Jan. 25, 2020, 5 pages.

Finkel et al., "Joint Parsing and Named Entity Recognition", Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the ACL, Jun. 2009, pp. 326-334.

Guo et al., "StateLens: A Reverse Engineering Solution for Making Existing Dynamic Touchscreens Accessible", In Proceedings of the 32nd Annual Symposium on User Interface Software and Technology (UIST '19), New Orleans, LA, USA, Online available at: <https://dl.acm.org/doi/pdf/10.1145/3332165.3347873>, Oct. 20-23, 2019, pp. 371-385.

Gupta et al., "I-vector-based Speaker Adaptation of Deep Neural Networks for French Broadcast Audio Transcription", ICASSP, 2014, 2014, pp. 6334-6338.

Heller et al., "AudioScope: Smartphones as Directional Microphones in Mobile Audio Augmented Reality Systems", In Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15), Crossings, Seoul, Korea, Online available at: <https://dl.acm.org/doi/pdf/10.1145/2702123.2702159>, Apr. 18-23, 2015, pp. 949-952.

Lin Luyuan, "An Assistive Handwashing System with Emotional Intelligence", Using Emotional Intelligence in Cognitive Intelligent Assistant Systems, 2014, 101 pages.

Muller et al., "Control Theoretic Models of Pointing", ACM Transactions on Computer-Human Interaction, Aug. 2017, 36 pages.

Pan et al., "Natural Language Aided Visual Query Building for Complex Data Access", In proceeding of: Proceedings of the Twenty-Second Conference on Innovative Applications of Artificial Intelligence, XP055114607, Jul. 11, 2010, pp. 1821-1826.

Phoenix Solutions, Inc., "Declaration of Christopher Schmandt Regarding the MIT Galaxy System", West Interactive Corp., a Delaware Corporation, Document 40, Jul. 2, 2010, 162 pages.

Rodrigues et al., "Exploring Mixed Reality in Specialized Surgical Environments", In Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '17), Denver, CO, USA, Online available at: <https://dl.acm.org/doi/pdf/10.1145/3027063.3053273>, May 6-11, 2017, pp. 2591-2598.

Rowland et al., "Designing Connected Products: UX for the Consumer Internet of Things", O'Reilly, May 31, 2015, 452 pages.

Senior et al., "Improving DNN Speaker Independence With I-Vector Inputs", ICASSP, 2014, pp. 225-229.

Seroter et al., "SOA Patterns with BizTalk Server 2013 and Microsoft Azure", Second Edition, Packt Publishing, Jun. 30, 2015, 454 pages.

Xu et al., "Speech-Based Interactive Games for Language Learning: Reading, Translation, and Question-Answering", Computational Linguistics and Chinese Language Processing, vol. 14, No. 2, Jun. 2009, pp. 133-160.

* cited by examiner

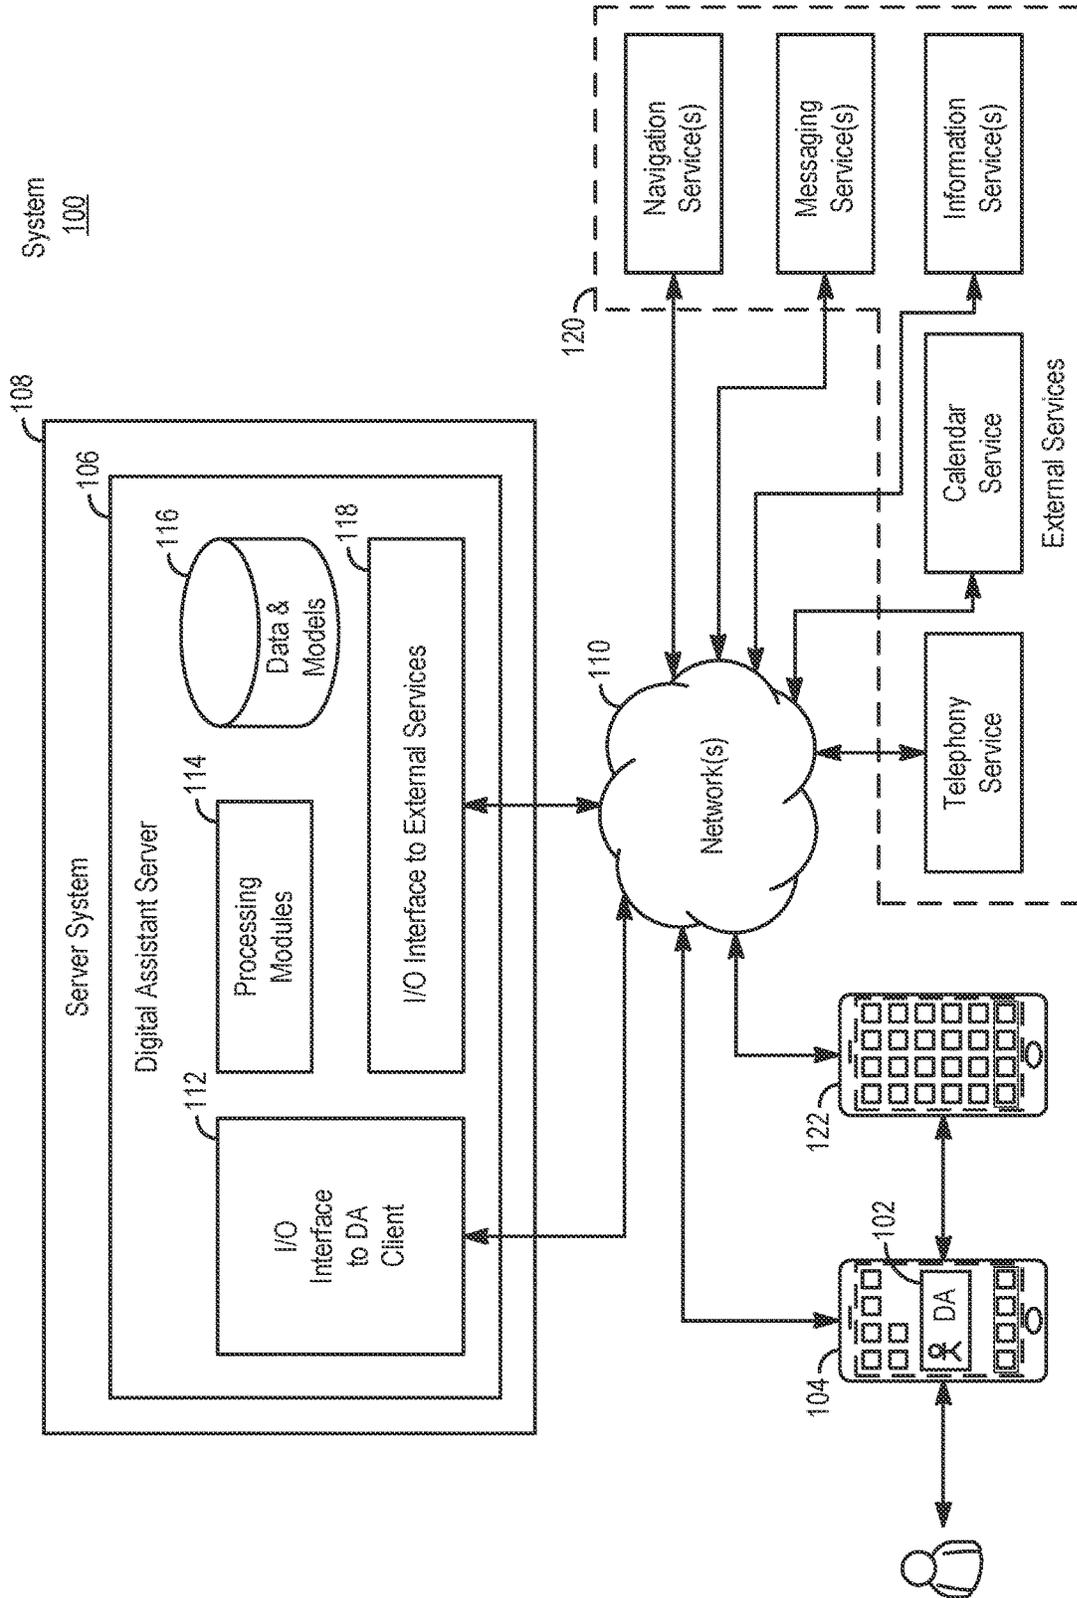


FIG. 1

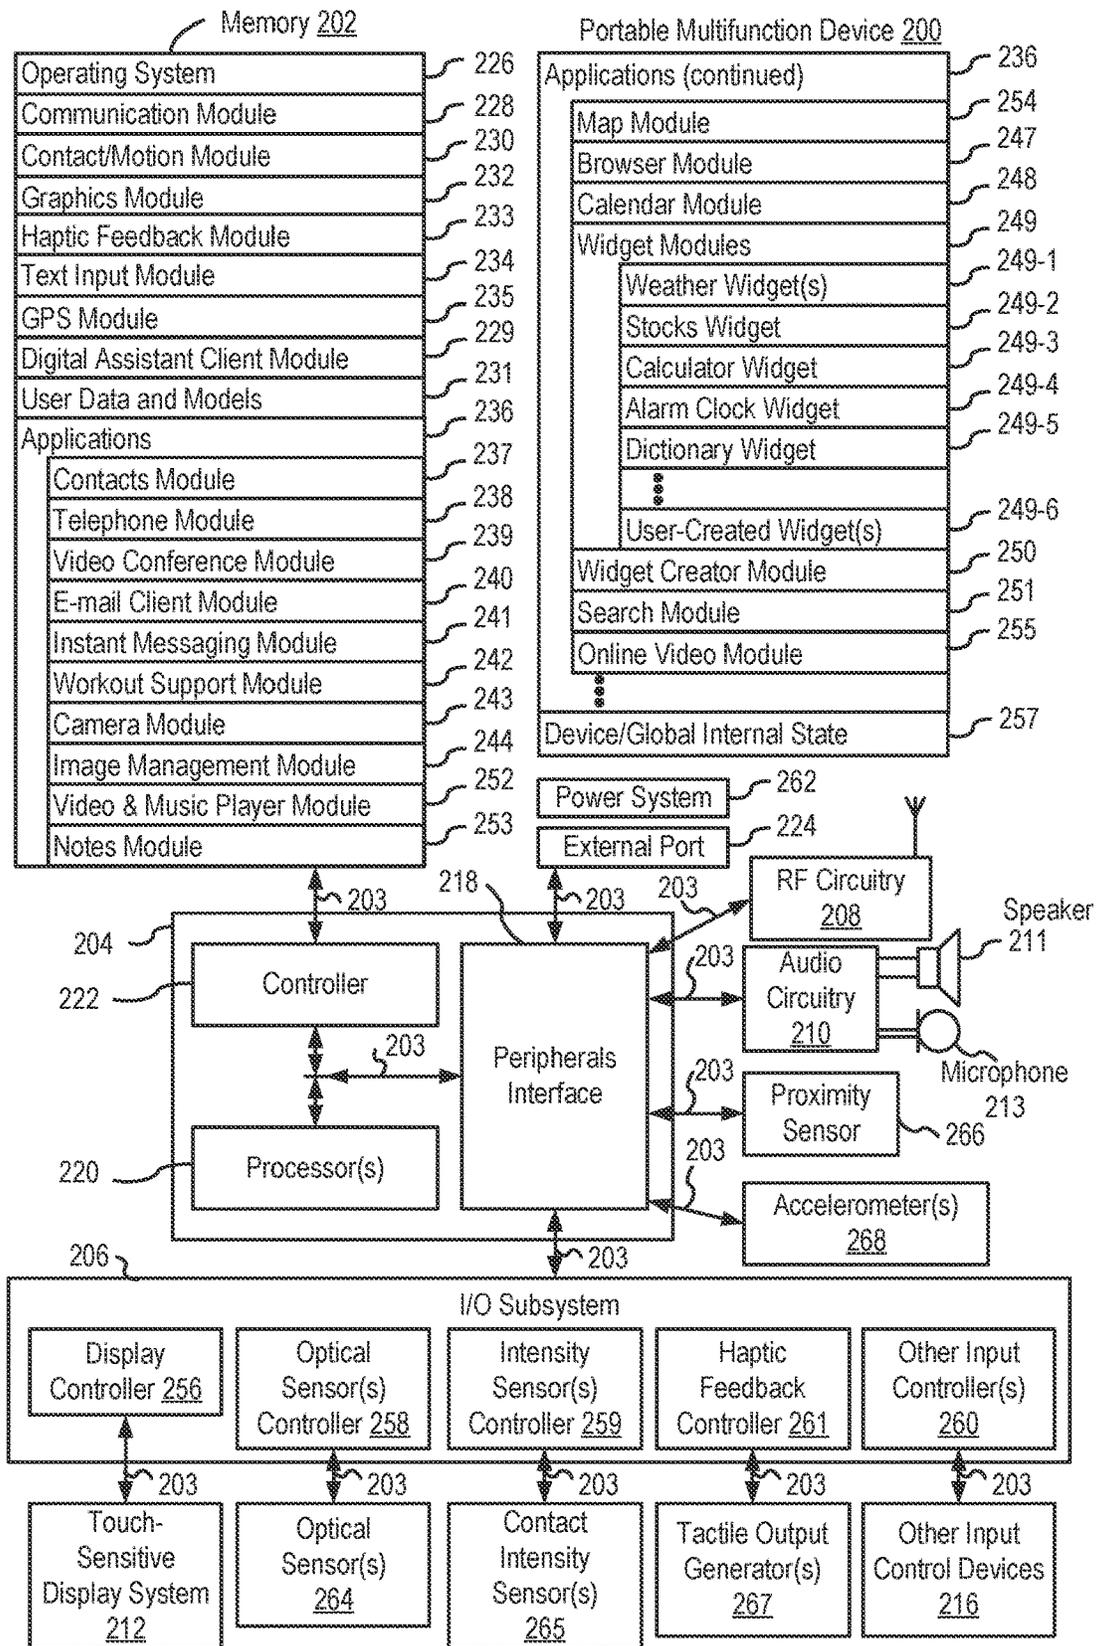


FIG. 2A

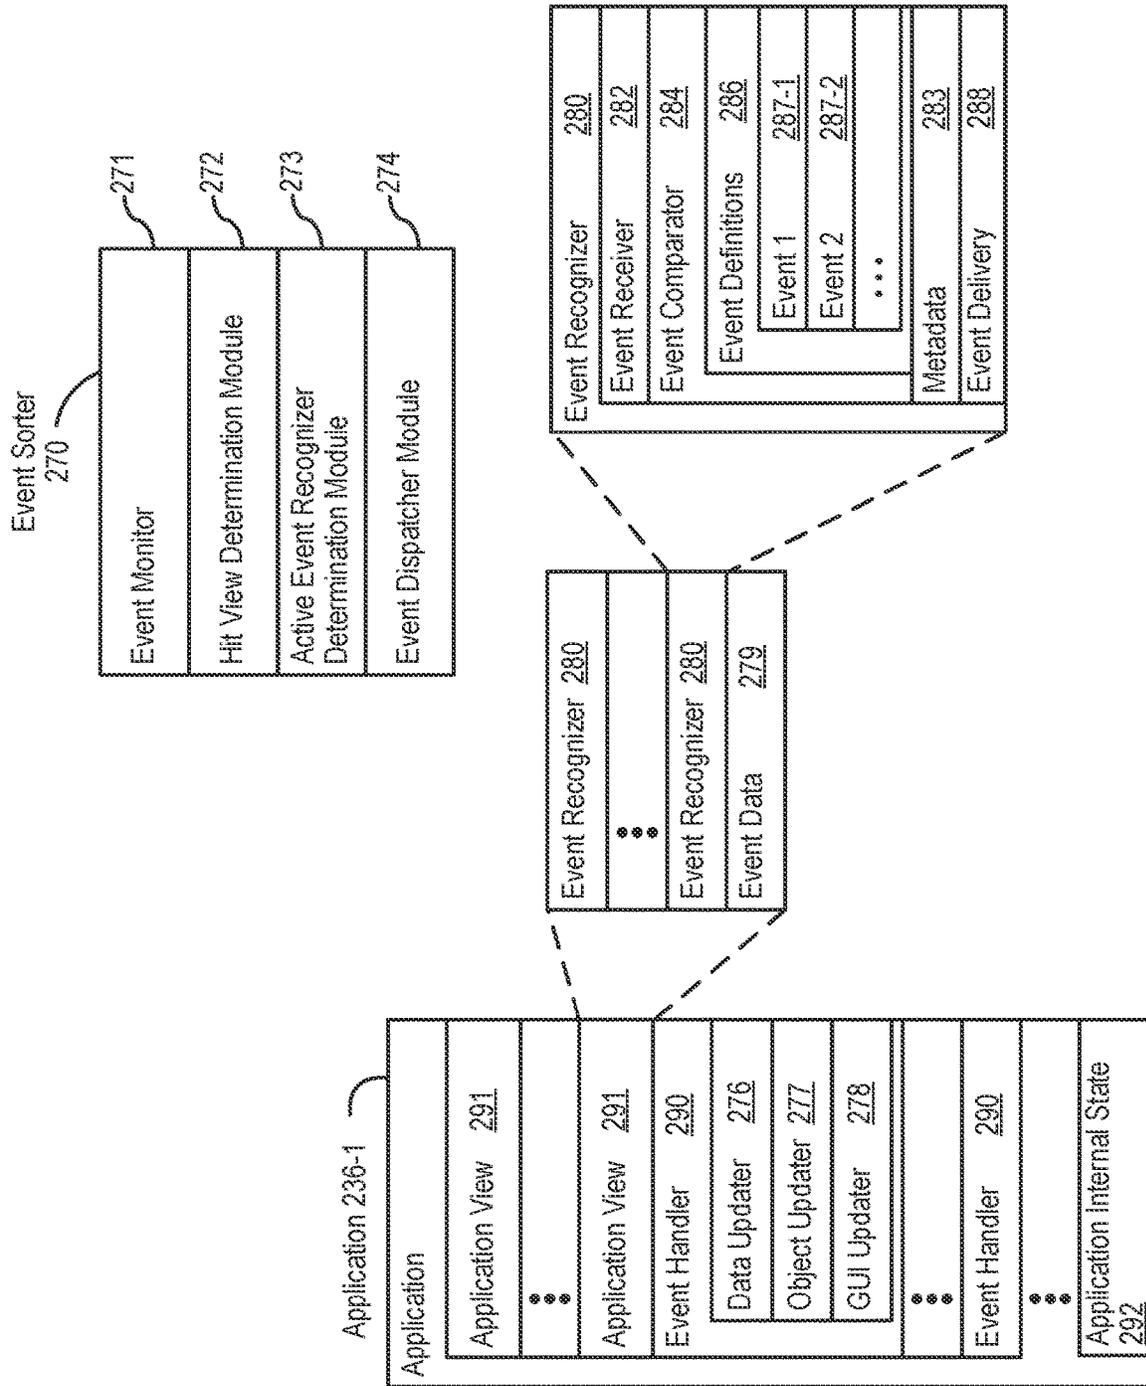


FIG. 2B

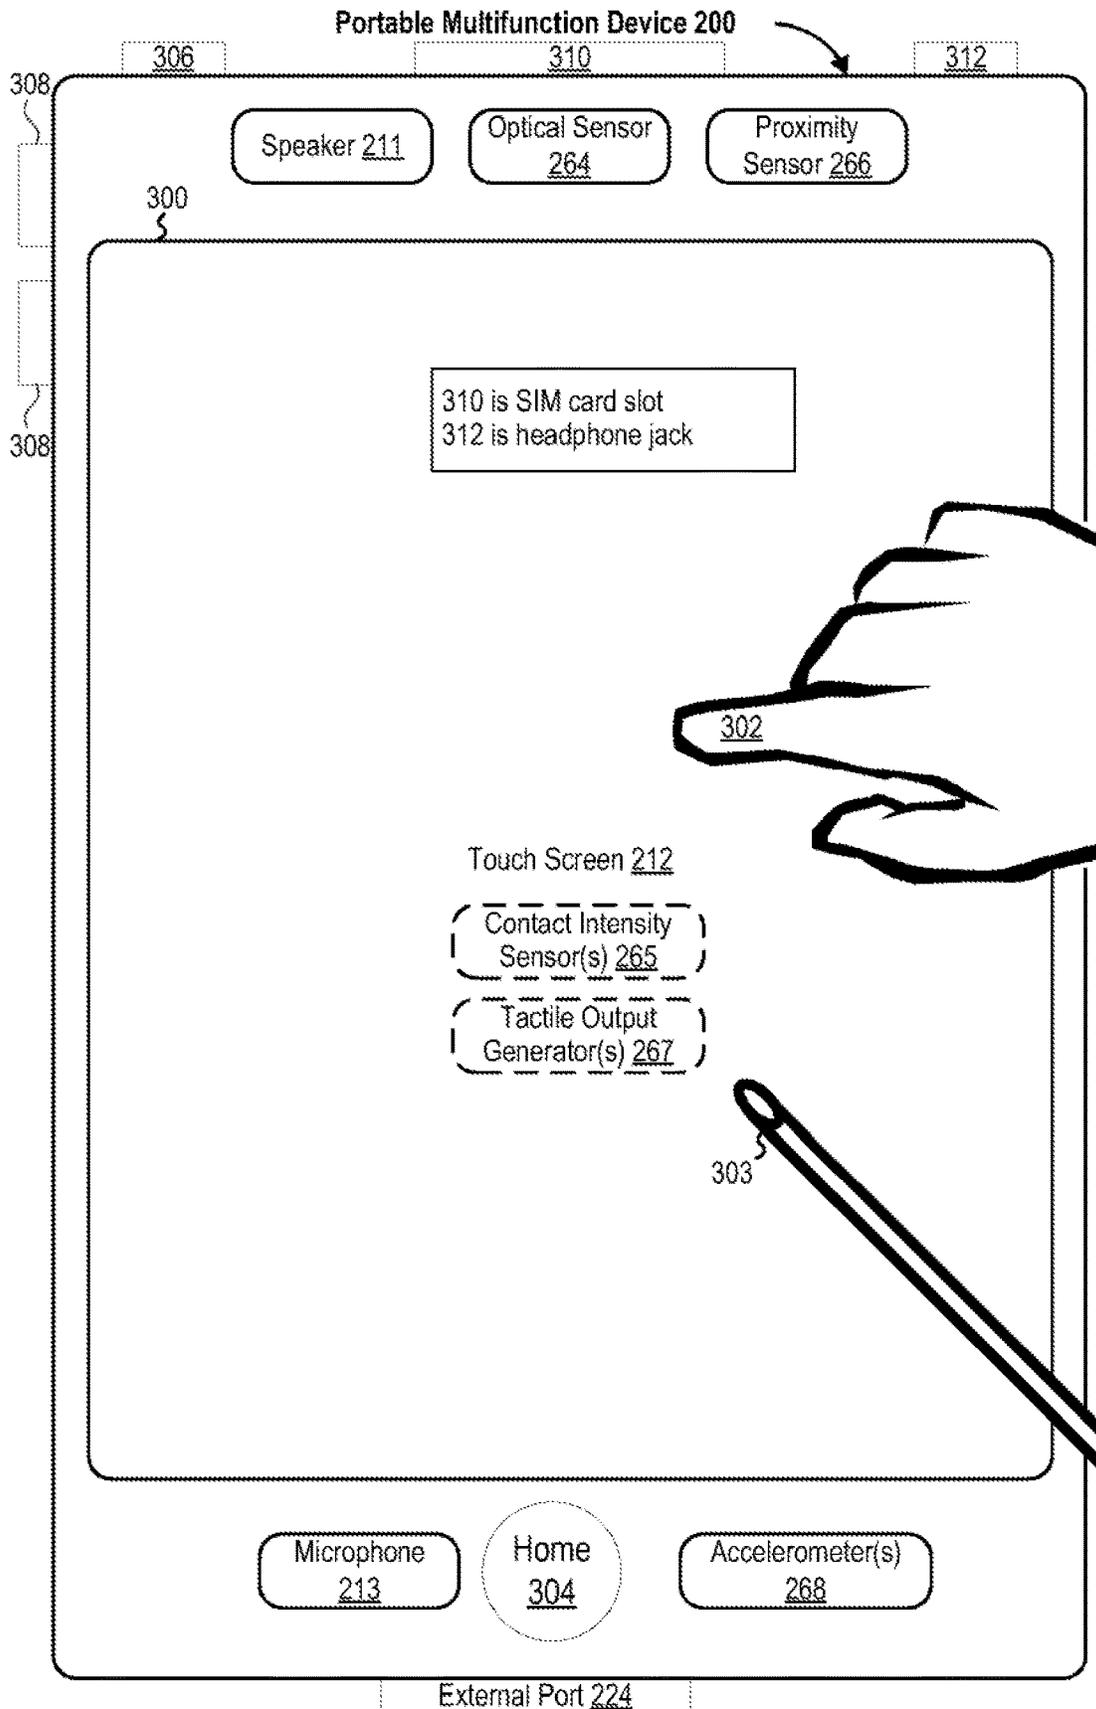


FIG. 3

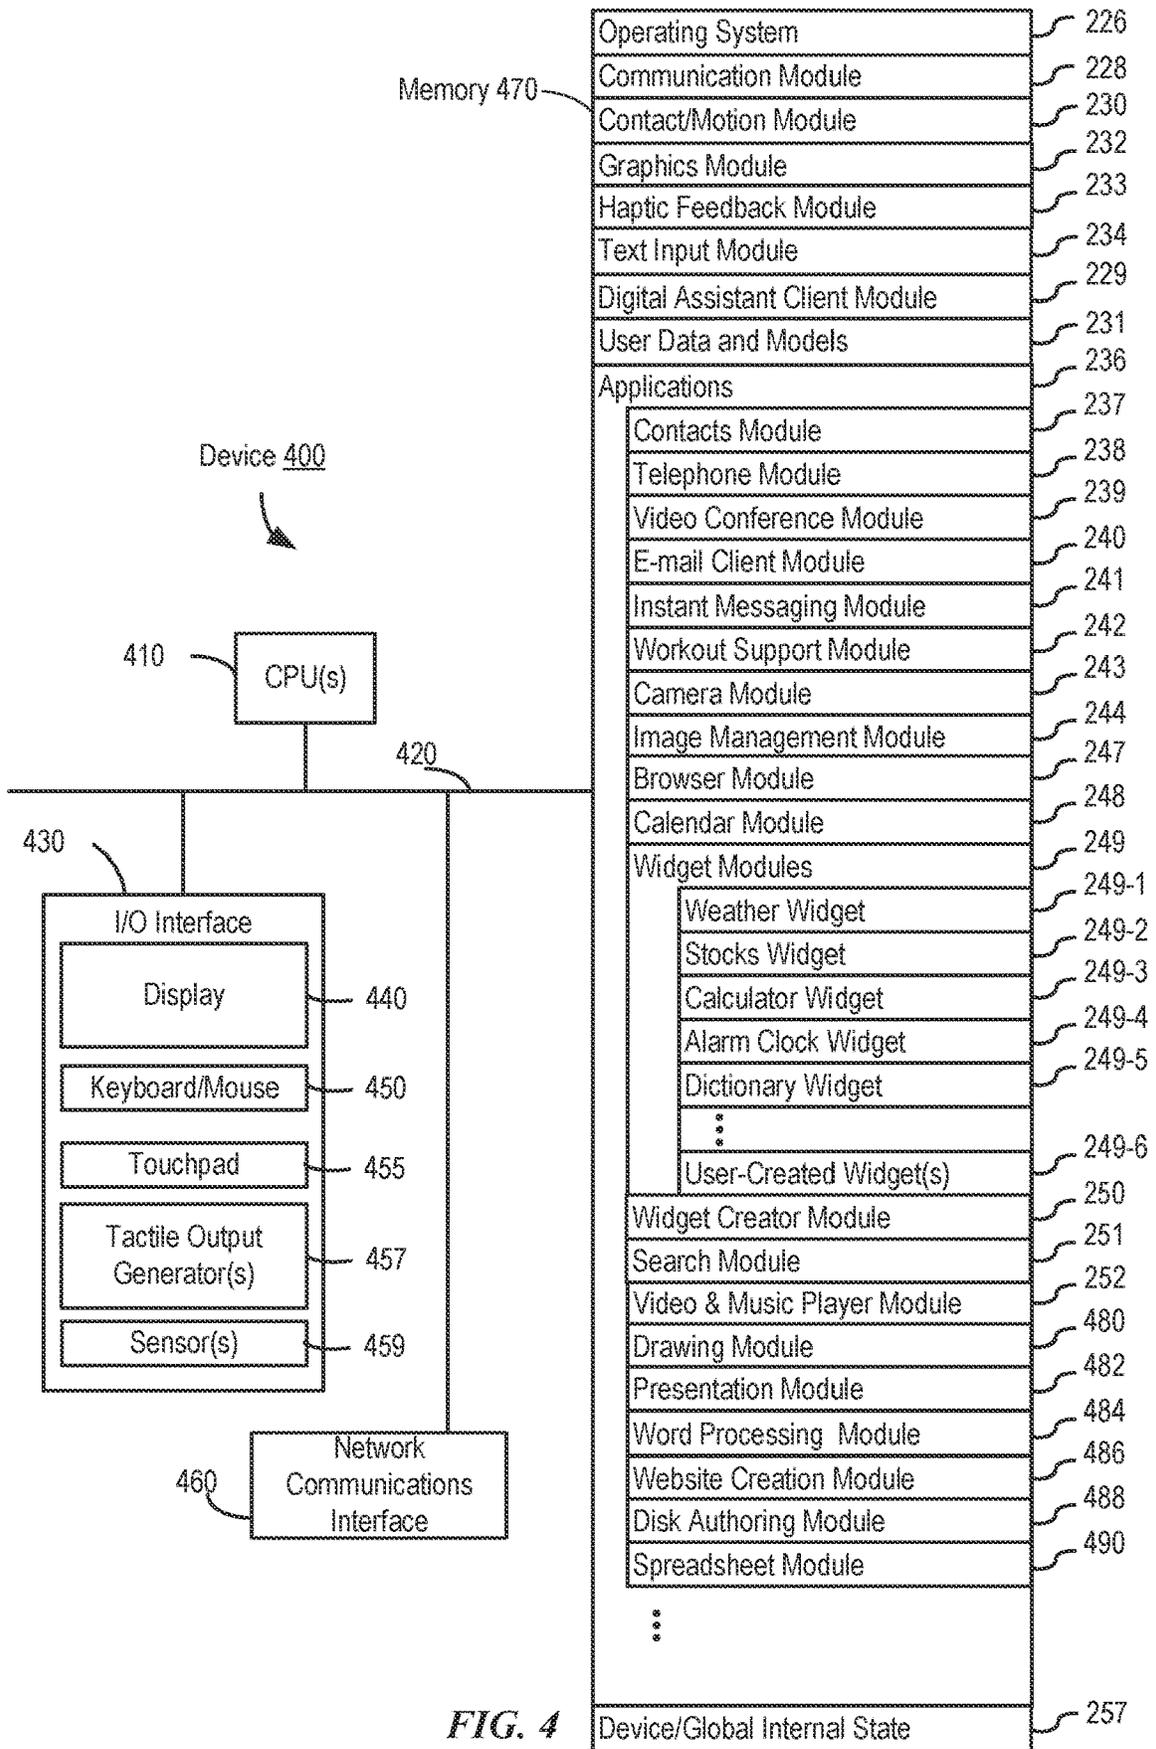


FIG. 4

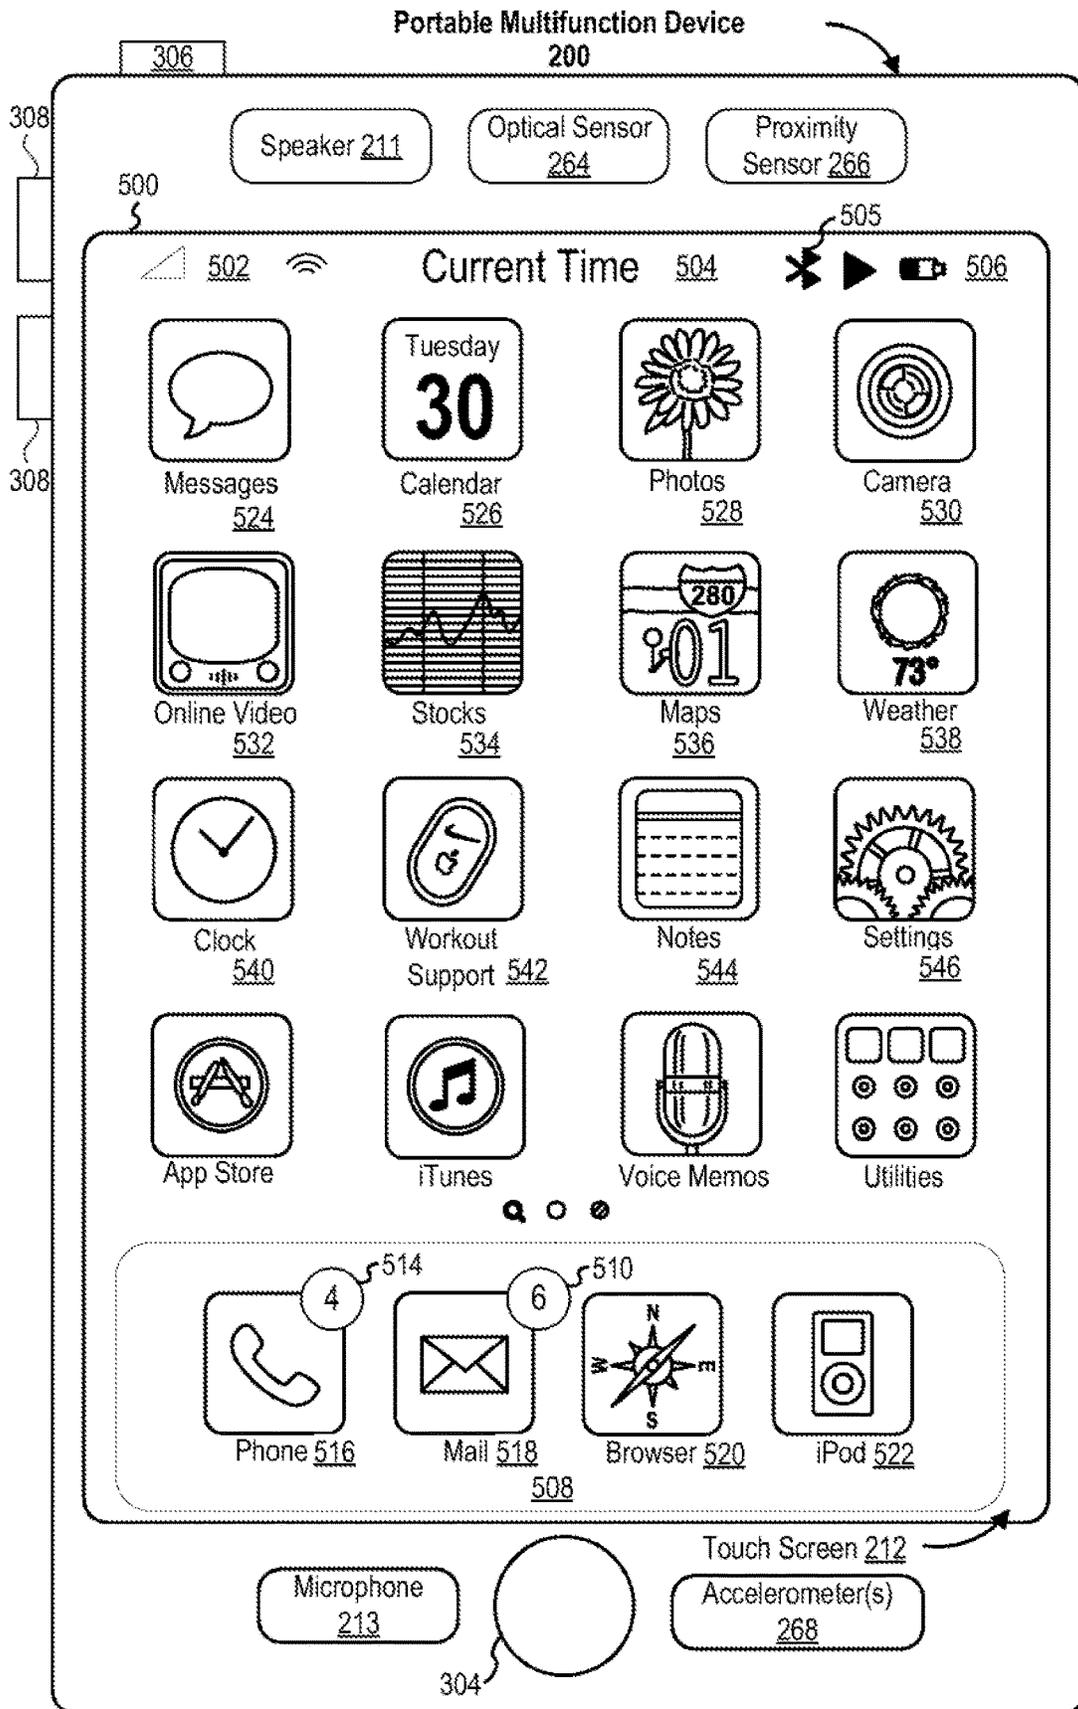


FIG. 5A

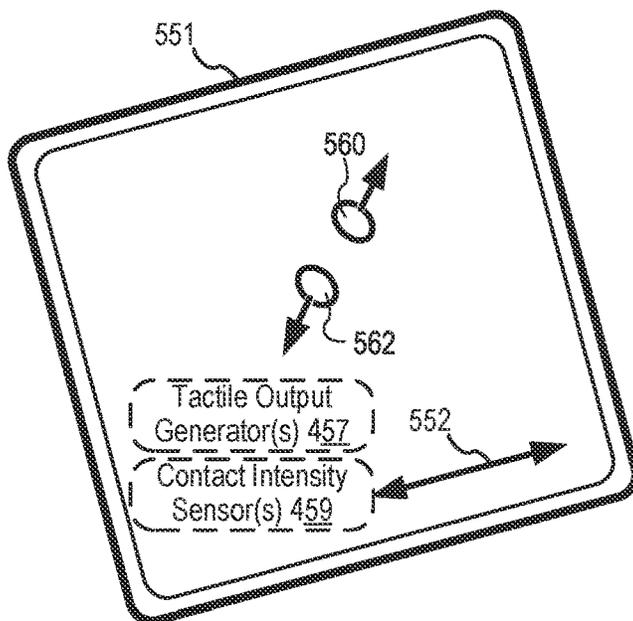
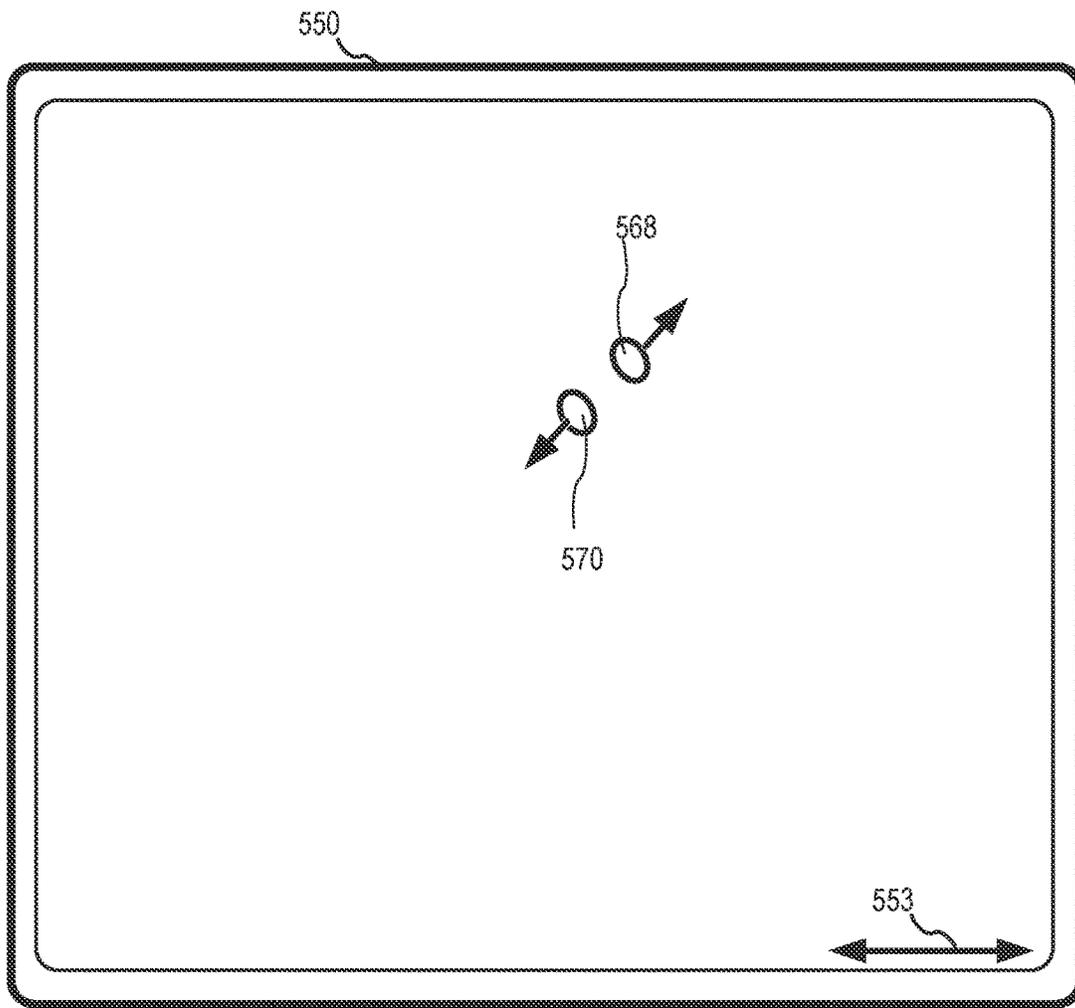


FIG. 5B

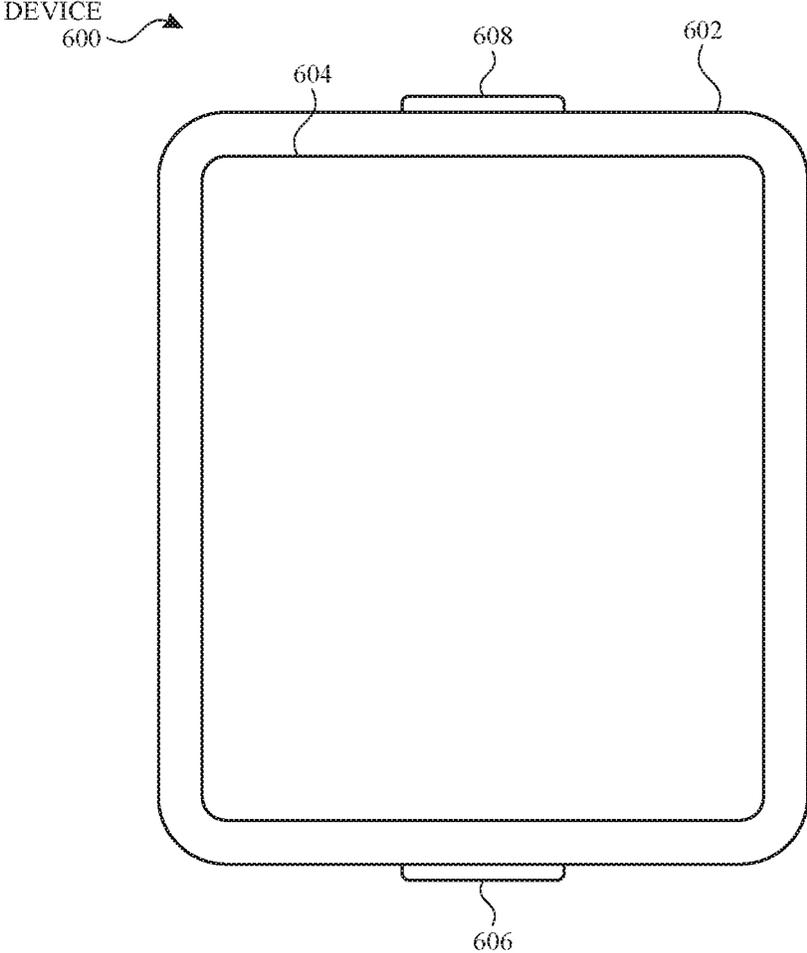


FIG. 6A

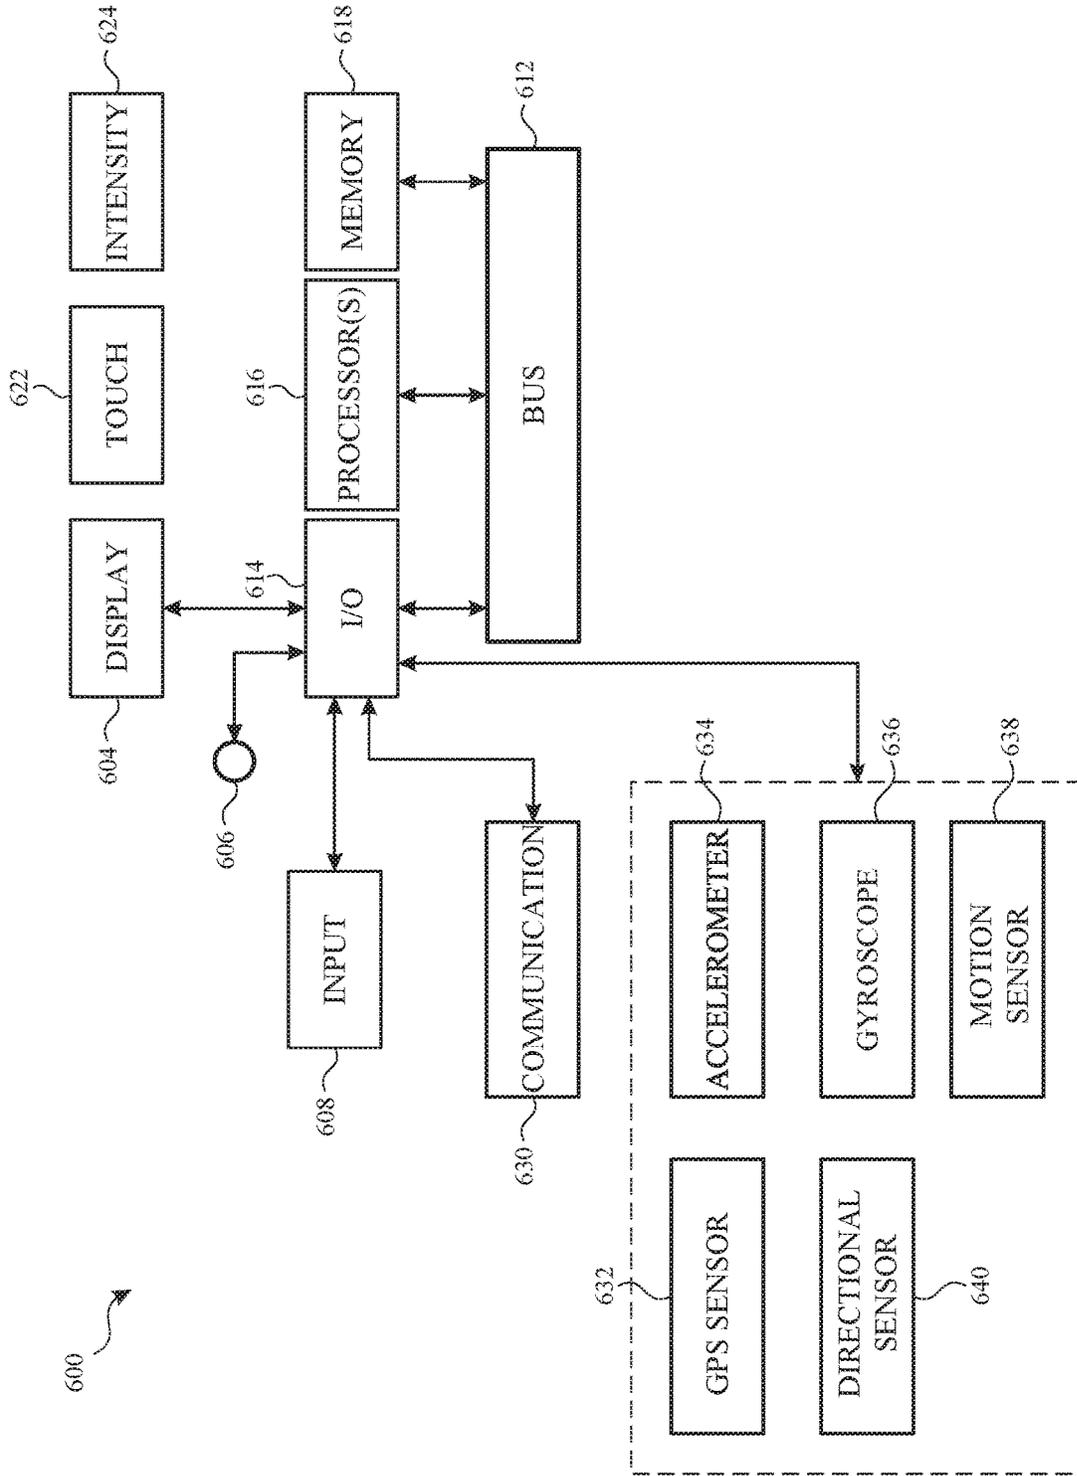


FIG. 6B

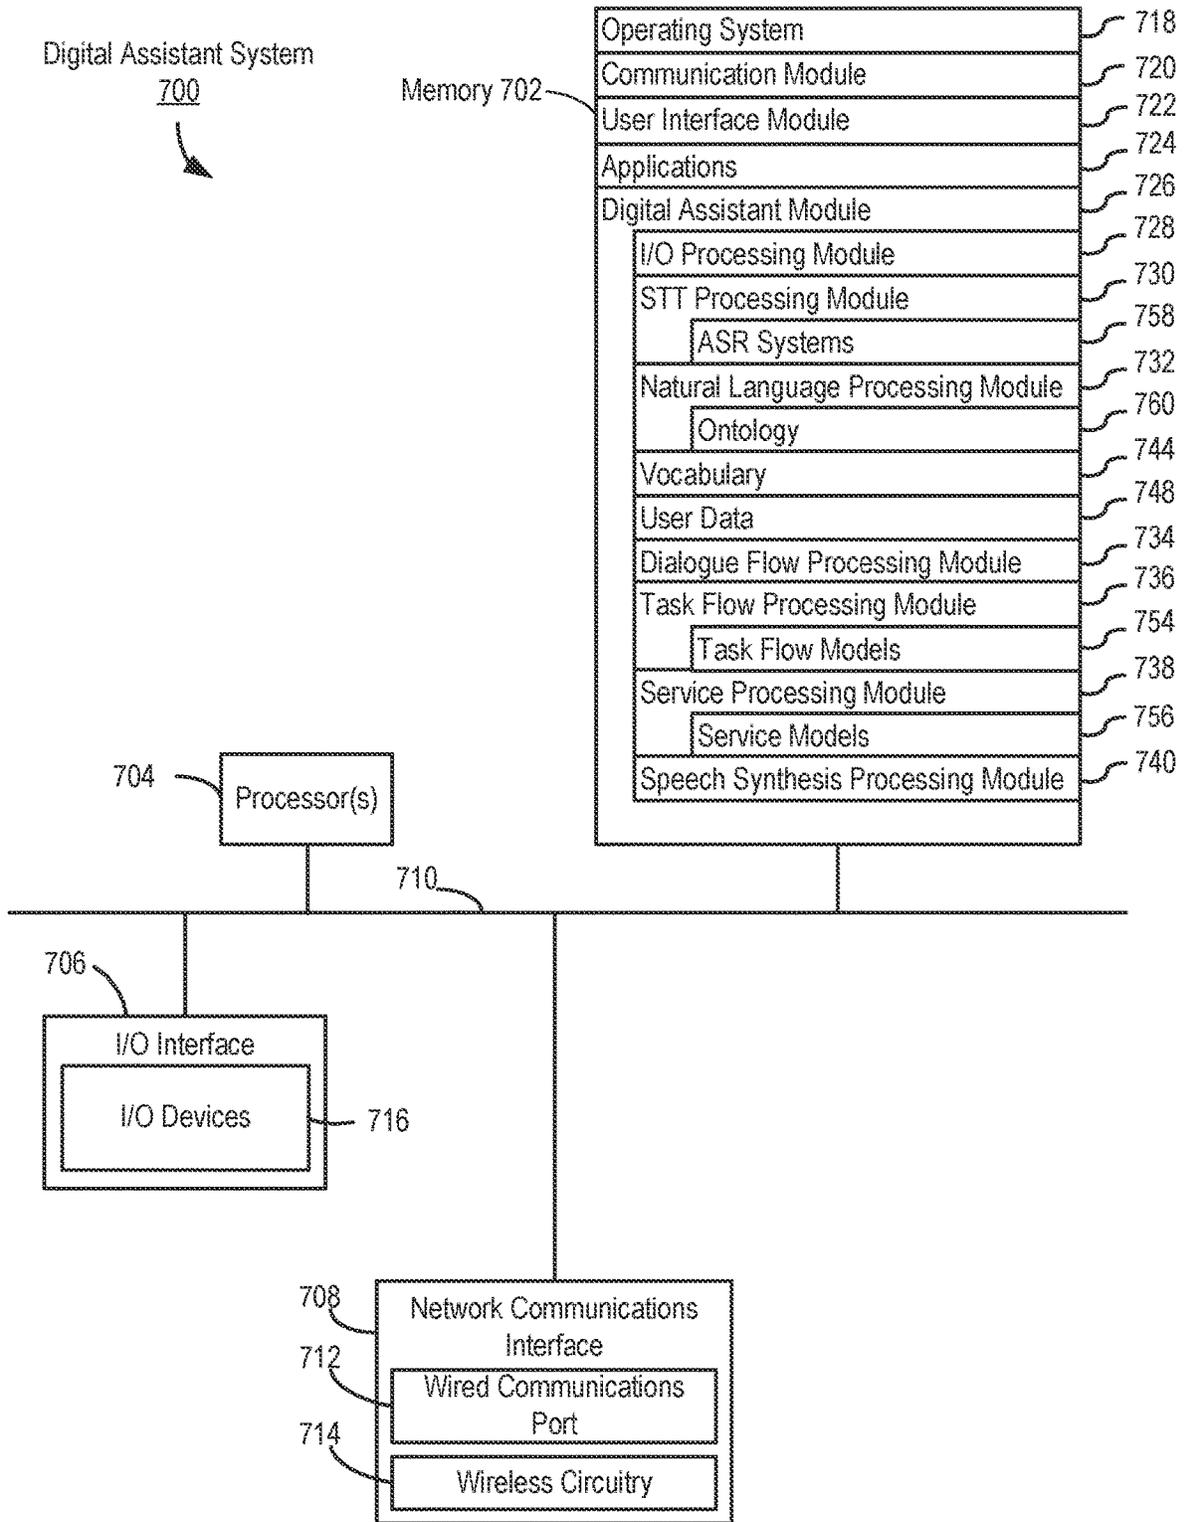


FIG. 7A

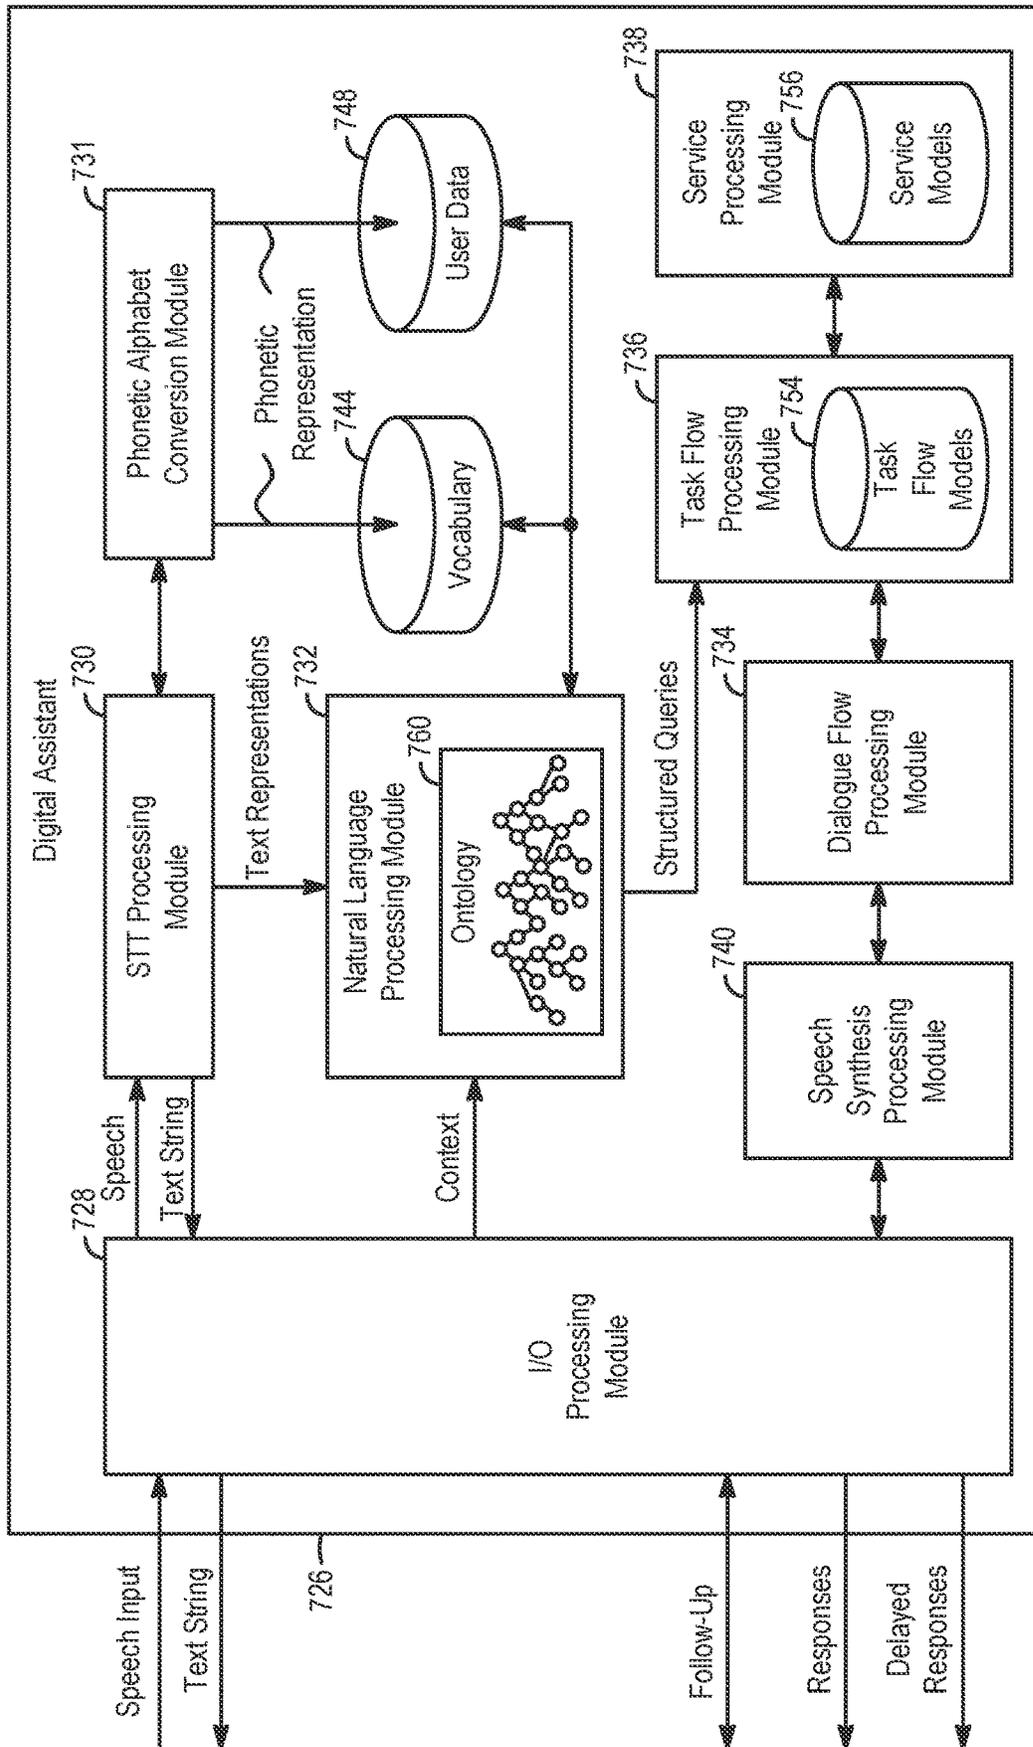


FIG. 7B

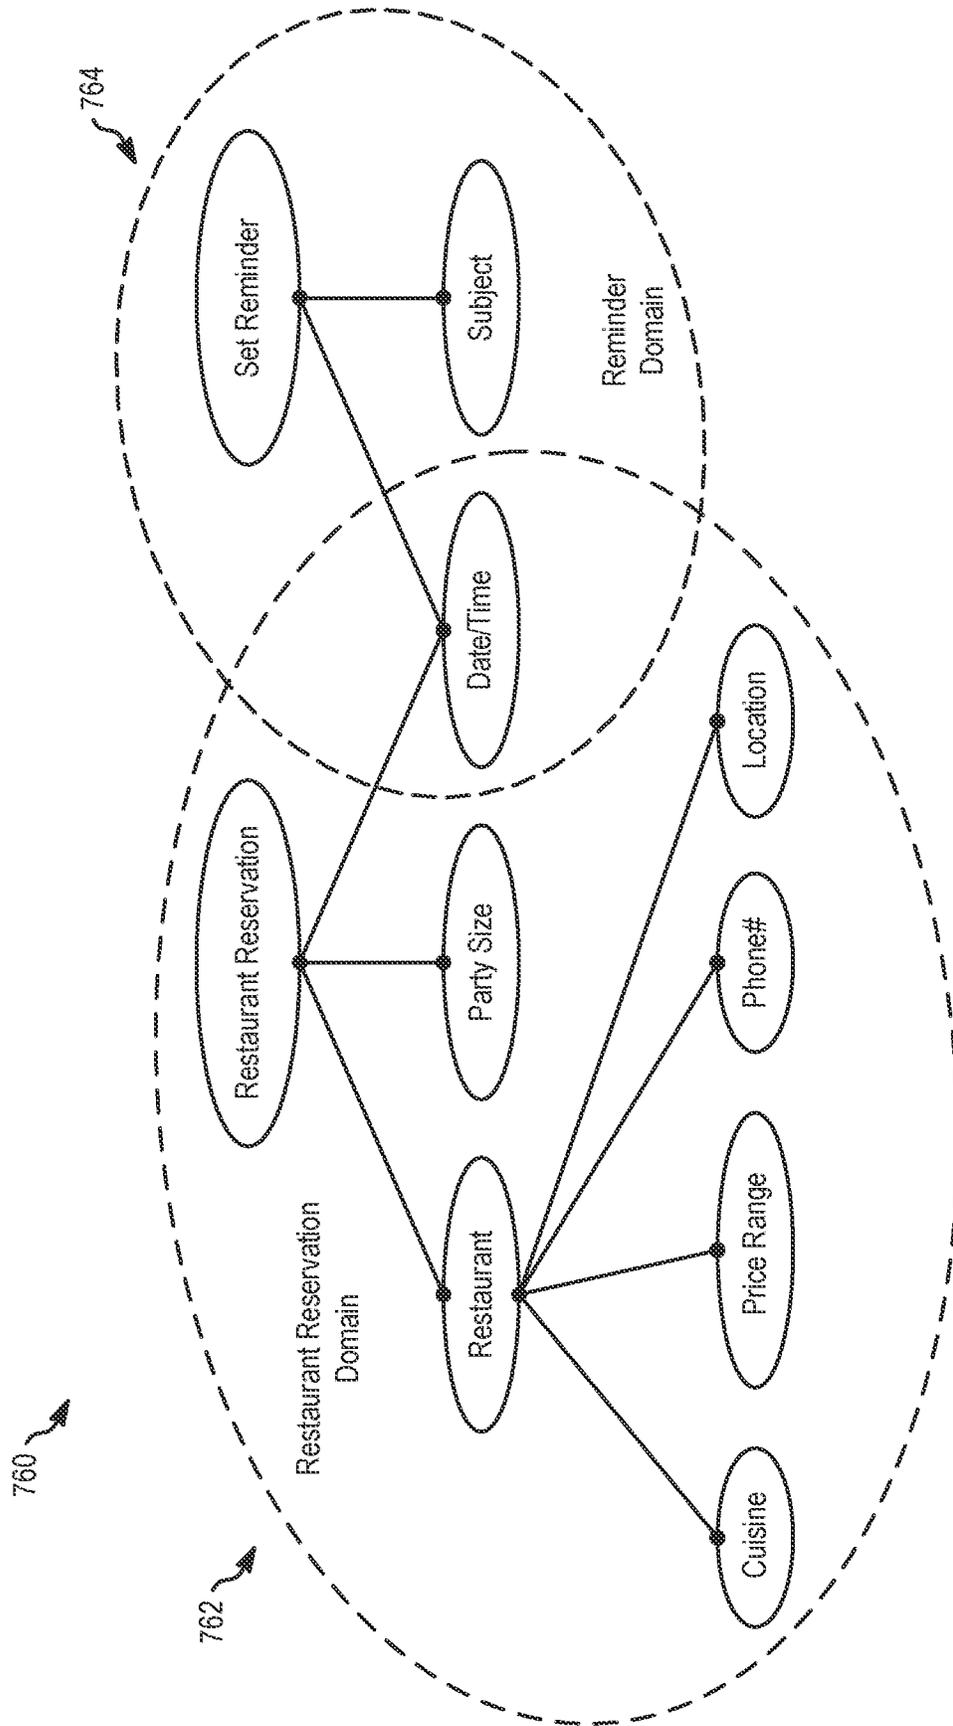


FIG. 7C

800

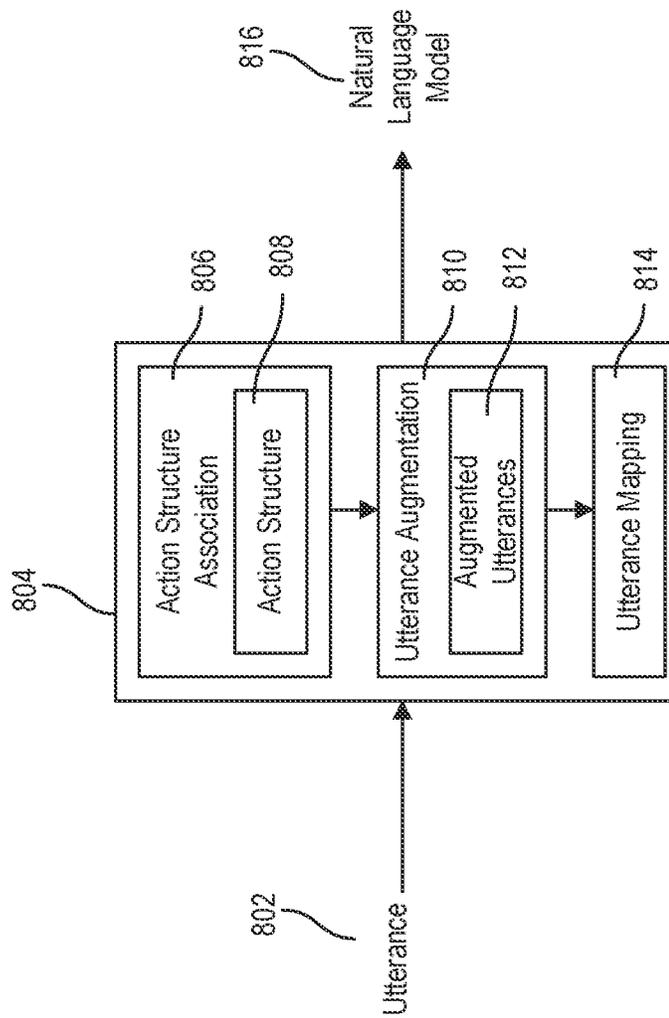


FIG. 8

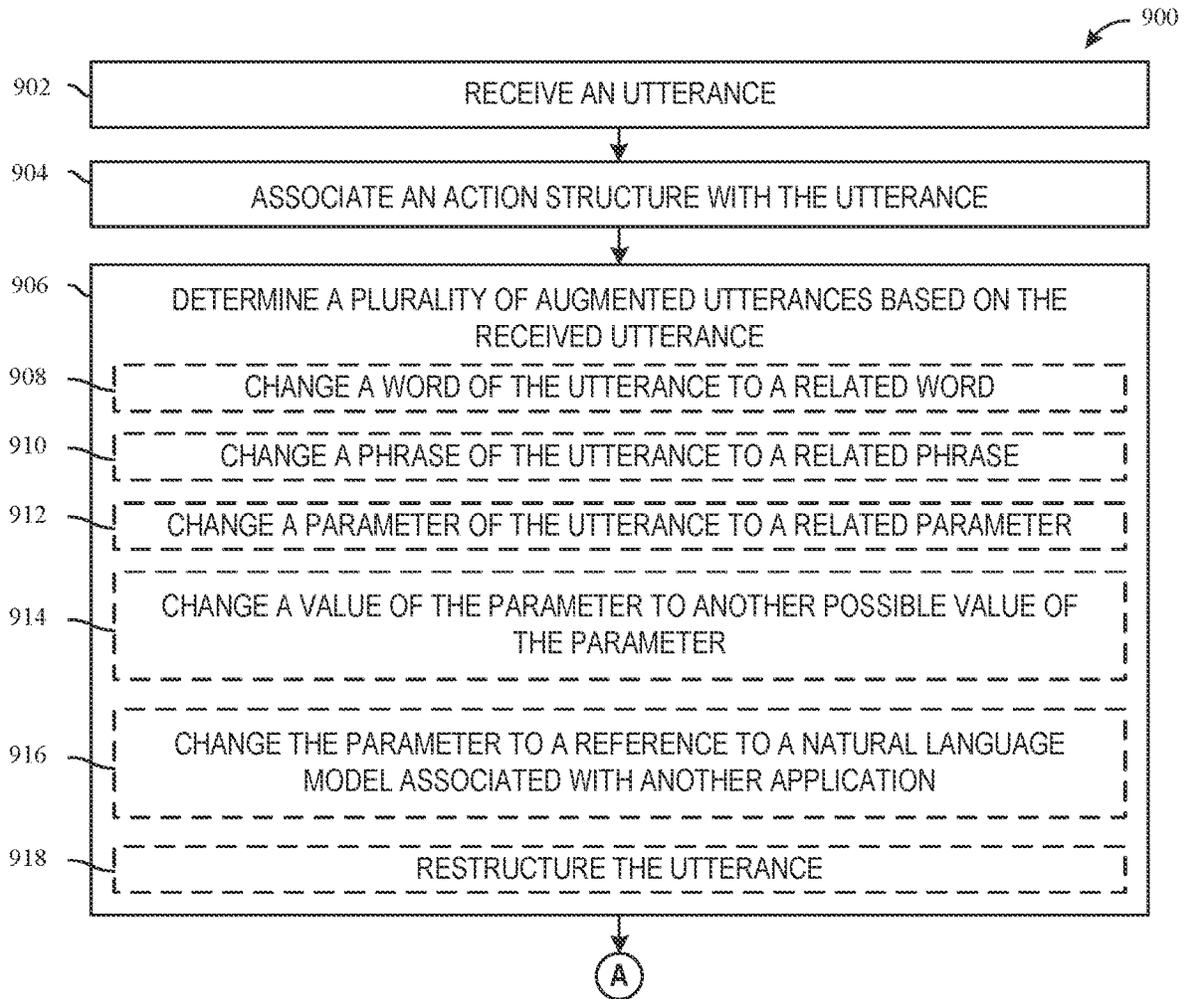


FIG. 9A

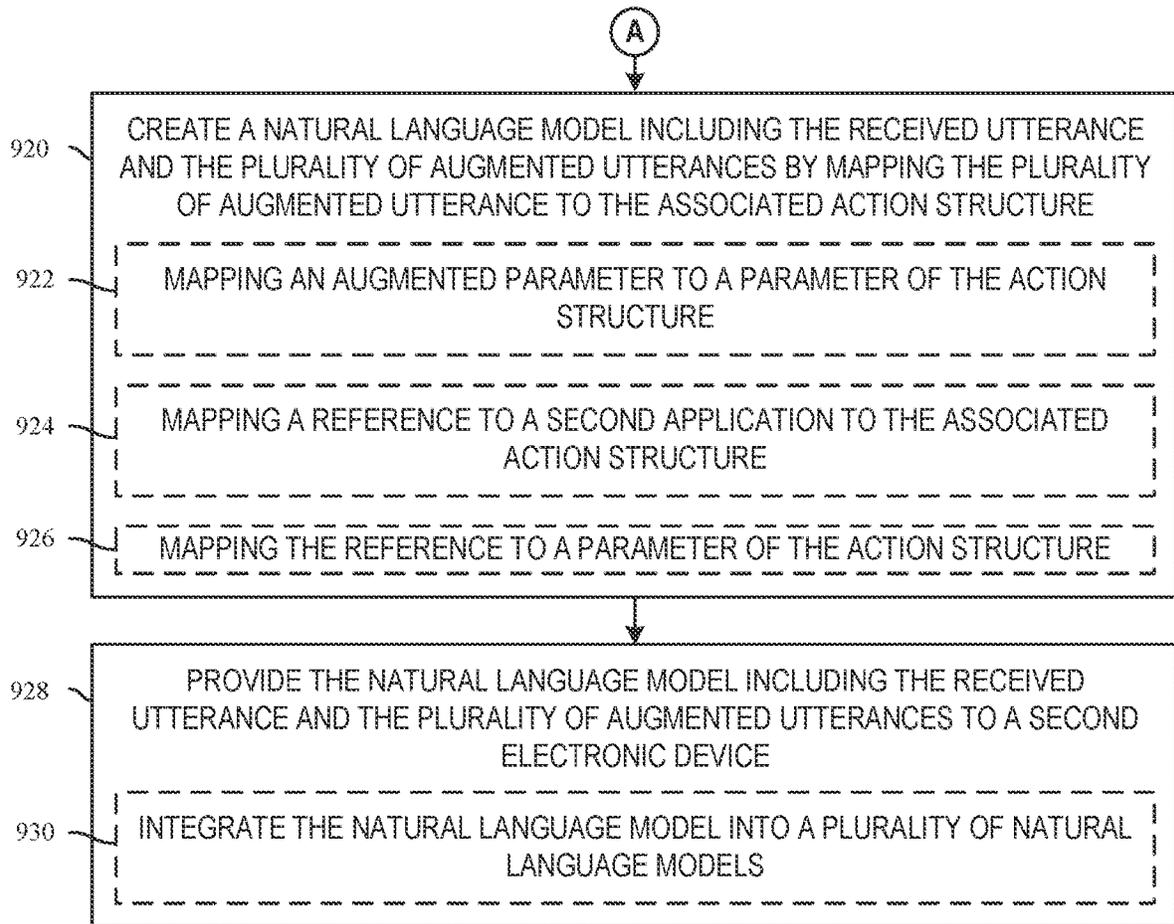


FIG. 9B

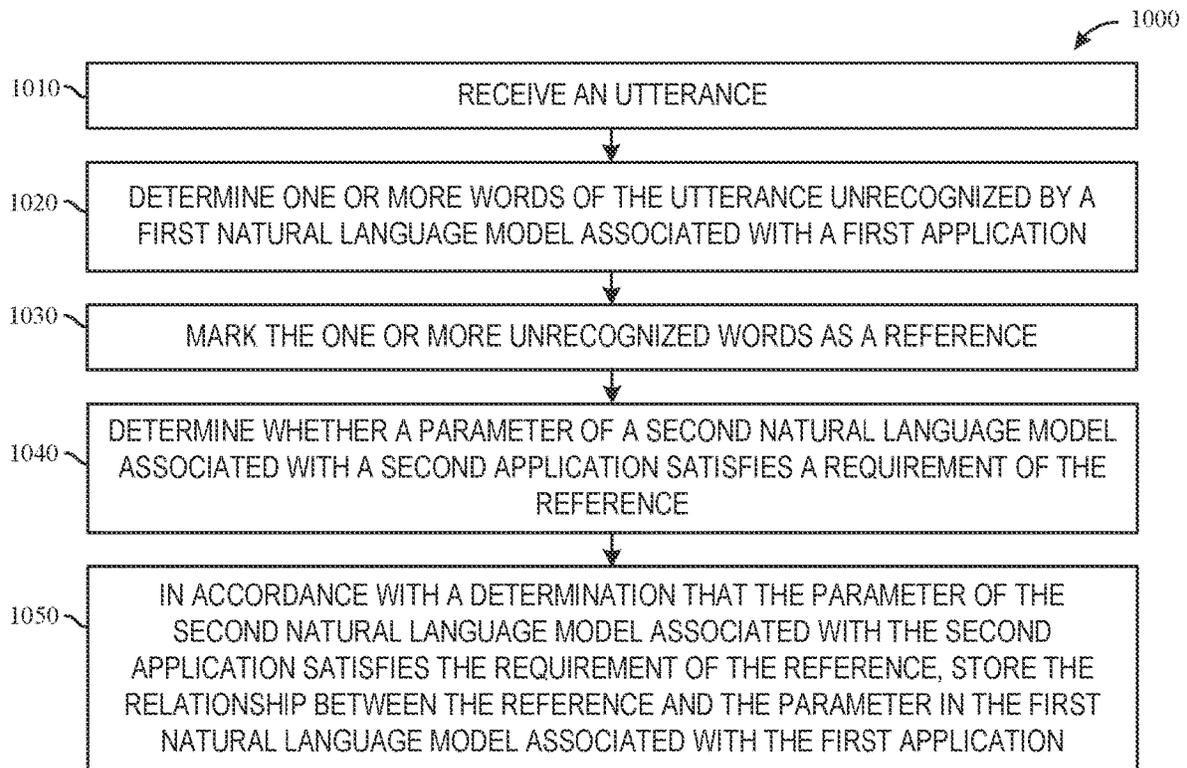


FIG. 10

1

USING AUGMENTATION TO CREATE NATURAL LANGUAGE MODELS

CROSS REFERENCE TO RELATED APPLICATIONS

This application claims the benefit of U.S. Provisional Application No. 63/078,280, entitled "USING AUGMENTATION TO CREATE NATURAL LANGUAGE MODELS," filed Sep. 14, 2020, the content of which is hereby incorporated by reference in its entirety for all purposes.

FIELD

This relates generally to intelligent automated assistants and, more specifically, to creating and updating natural language models used by intelligent automated assistants.

BACKGROUND

Intelligent automated assistants (or digital assistants) can provide a beneficial interface between human users and electronic devices. Such assistants can allow users to interact with devices or systems using natural language in spoken and/or text forms. For example, a user can provide a speech input containing a user request to a digital assistant operating on an electronic device. The digital assistant can interpret the user's intent from the speech input and operationalize the user's intent into tasks. The tasks can then be performed by executing one or more services of the electronic device, and a relevant output responsive to the user request can be returned to the user.

In some circumstances, interpreting the user intent may require natural language models that include a wide variety of possible utterances. For example, some natural language models may be able to recognize a limited number of utterances and determine a possible task from those utterances. Thus, when a new utterance is received that the natural language model has not been trained to recognize, the natural language model is unable to determine a task. Accordingly, it is advantageous for natural language models to be generated based on a wider variety of utterances without requiring substantially more work from developers or creators of the natural language models.

SUMMARY

Example methods are disclosed herein. An example method includes, at an electronic device with one or more processors and memory, receiving an utterance, associating an action structure with the utterance, determining a plurality of augmented utterances based on the received utterance, creating a natural language model including the received utterance and the plurality of augmented utterances by mapping the plurality of augmented utterance to the associated action structure, and providing the natural language model including the received utterance and the plurality of augmented utterances to a second electronic device.

Example non-transitory computer-readable media are disclosed herein. An example non-transitory computer-readable storage medium stores one or more programs. The one or more programs comprise instructions, which when executed by one or more processors of an electronic device, cause the electronic device to receive an utterance, associate an action structure with the utterance, determine a plurality of augmented utterances based on the received utterance, create a natural language model including the received utterance and

2

the plurality of augmented utterances by mapping the plurality of augmented utterance to the associated action structure, and provide the natural language model including the received utterance and the plurality of augmented utterances to a second electronic device.

Example electronic devices are disclosed herein. An example electronic device comprises one or more processors; a memory; and one or more programs, where the one or more programs are stored in the memory and configured to be executed by the one or more processors, the one or more programs including instructions for receiving an utterance, associating an action structure with the utterance, determining a plurality of augmented utterances based on the received utterance, creating a natural language model including the received utterance and the plurality of augmented utterances by mapping the plurality of augmented utterance to the associated action structure, and providing the natural language model including the received utterance and the plurality of augmented utterances to a second electronic device.

An example electronic device comprises means for receiving an utterance, associating an action structure with the utterance, determining a plurality of augmented utterances based on the received utterance, creating a natural language model including the received utterance and the plurality of augmented utterances by mapping the plurality of augmented utterance to the associated action structure, and providing the natural language model including the received utterance and the plurality of augmented utterances to a second electronic device.

Determining a plurality of augmented utterances based on the received utterance and creating a natural language model including the received utterance and the plurality of augmented utterances by mapping the plurality of augmented utterance to the associated action structure allows for more efficient and automated creation of natural language models. For example, rather than requiring a developer of an application or other piece of software to manually provide the various inputs that a user may provide to the application and train the application to appropriately respond to each of the inputs, a smaller set of possible inputs may be provided and many different possible alternatives may be automatically generated. Creating the natural language model in this way is much more efficient than individually providing each possible speech input. Further, the natural language models created in this way cover a wider variety of possible spoken inputs and responses, including those that the developer or creator of the application may not be aware. Thus, creating natural language models in this way increases the efficiency and the responsiveness of a digital assistant and other applications installed on a device which may be the focus of a user request or input.

An example method includes, at an electronic device with one or more processors and memory, receiving an utterance, determining one or more words of the utterance unrecognized by a first natural language model associated with a first application, marking the one or more unrecognized words as a reference, determining whether a parameter of a second natural language model associated with a second application satisfies a requirement of the reference, and in accordance with a determination that the parameter of the second natural language model associated with the second application satisfies the requirement of the reference, storing the relationship between the reference and the parameter in the first natural language model associated with the first application.

An example non-transitory computer-readable storage medium stores one or more programs. The one or more

programs comprise instructions, which when executed by one or more processors of an electronic device, cause the electronic device to receive an utterance, determine one or more words of the utterance unrecognized by a first natural language model associated with a first application, mark the one or more unrecognized words as a reference, determine whether a parameter of a second natural language model associated with a second application satisfies a requirement of the reference, and in accordance with a determination that the parameter of the second natural language model associated with the second application satisfies the requirement of the reference, store the relationship between the reference and the parameter in the first natural language model associated with the first application.

An example electronic device comprises one or more processors; a memory; and one or more programs, where the one or more programs are stored in the memory and configured to be executed by the one or more processors, the one or more programs including instructions for receiving an utterance, determining one or more words of the utterance unrecognized by a first natural language model associated with a first application, marking the one or more unrecognized words as a reference, determining whether a parameter of a second natural language model associated with a second application satisfies a requirement of the reference, and in accordance with a determination that the parameter of the second natural language model associated with the second application satisfies the requirement of the reference, storing the relationship between the reference and the parameter in the first natural language model associated with the first application.

An example electronic device comprises means for receiving an utterance, determining one or more words of the utterance unrecognized by a first natural language model associated with a first application, marking the one or more unrecognized words as a reference, determining whether a parameter of a second natural language model associated with a second application satisfies a requirement of the reference, and in accordance with a determination that the parameter of the second natural language model associated with the second application satisfies the requirement of the reference, storing the relationship between the reference and the parameter in the first natural language model associated with the first application.

Marking the one or more unrecognized words as a reference and determining whether a parameter of a second natural language model associated with a second application satisfies a requirement of the reference allows for more efficient and responsive processing of natural language inputs. In particular, by determining and processing unrecognized words in this manner the natural language models are automatically updated to interact and reference other natural language models, increasing the interconnectivity of the available models. In this way the natural language models may adapt overtime to respond to a wider variety of inputs in a more efficient manner, reducing the processing required and increasing user satisfaction by providing the correct task determination from the natural language input efficiently.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram illustrating a system and environment for implementing a digital assistant, according to various examples.

FIG. 2A is a block diagram illustrating a portable multifunction device implementing the client-side portion of a digital assistant, according to various examples.

FIG. 2B is a block diagram illustrating exemplary components for event handling, according to various examples.

FIG. 3 illustrates a portable multifunction device implementing the client-side portion of a digital assistant, according to various examples.

FIG. 4 is a block diagram of an exemplary multifunction device with a display and a touch-sensitive surface, according to various examples.

FIG. 5A illustrates an exemplary user interface for a menu of applications on a portable multifunction device, according to various examples.

FIG. 5B illustrates an exemplary user interface for a multifunction device with a touch-sensitive surface that is separate from the display, according to various examples.

FIG. 6A illustrates a personal electronic device, according to various examples.

FIG. 6B is a block diagram illustrating a personal electronic device, according to various examples.

FIG. 7A is a block diagram illustrating a digital assistant system or a server portion thereof, according to various examples.

FIG. 7B illustrates the functions of the digital assistant shown in FIG. 7A, according to various examples.

FIG. 7C illustrates a portion of an ontology, according to various examples.

FIG. 8 illustrates a system for creating a natural language model, according to various examples.

FIGS. 9A-9B illustrates a process for creating a natural language model, according to various examples.

FIG. 10 illustrates a process for updating a natural language model, according to various examples.

DETAILED DESCRIPTION

In the following description of examples, reference is made to the accompanying drawings in which are shown by way of illustration specific examples that can be practiced. It is to be understood that other examples can be used and structural changes can be made without departing from the scope of the various examples.

In some examples, a method includes, at an electronic device with one or more processors and memory, receiving an utterance, associating an action structure with the utterance, determining a plurality of augmented utterances based on the received utterance, creating a natural language model including the received utterance and the plurality of augmented utterances by mapping the plurality of augmented utterance to the associated action structure, and providing the natural language model including the received utterance and the plurality of augmented utterances to an application. As further discussed below, creating natural language models in this way results in a more responsive digital assistant and applications that more freely interact with other applications and the digital assistant. Additionally, creating natural language models in this way is more efficient than previous methods of creating natural language models.

Although the following description uses terms “first,” “second,” etc. to describe various elements, these elements should not be limited by the terms. These terms are only used to distinguish one element from another. For example, a first input could be termed a second input, and, similarly, a second input could be termed a first input, without departing from the scope of the various described examples. The

first input and the second input are both inputs and, in some cases, are separate and different inputs.

The terminology used in the description of the various described examples herein is for the purpose of describing particular examples only and is not intended to be limiting. As used in the description of the various described examples and the appended claims, the singular forms “a,” “an,” and “the” are intended to include the plural forms as well, unless the context clearly indicates otherwise. It will also be understood that the term “and/or” as used herein refers to and encompasses any and all possible combinations of one or more of the associated listed items. It will be further understood that the terms “includes,” “including,” “comprises,” and/or “comprising,” when used in this specification, specify the presence of stated features, integers, steps, operations, elements, and/or components, but do not preclude the presence or addition of one or more other features, integers, steps, operations, elements, components, and/or groups thereof.

The term “if” may be construed to mean “when” or “upon” or “in response to determining” or “in response to detecting,” depending on the context. Similarly, the phrase “if it is determined” or “if [a stated condition or event] is detected” may be construed to mean “upon determining” or “in response to determining” or “upon detecting [the stated condition or event]” or “in response to detecting [the stated condition or event],” depending on the context.

1. System and Environment

FIG. 1 illustrates a block diagram of system 100 according to various examples. In some examples, system 100 implements a digital assistant. The terms “digital assistant,” “virtual assistant,” “intelligent automated assistant,” or “automatic digital assistant” refer to any information processing system that interprets natural language input in spoken and/or textual form to infer user intent, and performs actions based on the inferred user intent. For example, to act on an inferred user intent, the system performs one or more of the following: identifying a task flow with steps and parameters designed to accomplish the inferred user intent, inputting specific requirements from the inferred user intent into the task flow; executing the task flow by invoking programs, methods, services, APIs, or the like; and generating output responses to the user in an audible (e.g., speech) and/or visual form.

Specifically, a digital assistant is capable of accepting a user request at least partially in the form of a natural language command, request, statement, narrative, and/or inquiry. Typically, the user request seeks either an informational answer or performance of a task by the digital assistant. A satisfactory response to the user request includes a provision of the requested informational answer, a performance of the requested task, or a combination of the two. For example, a user asks the digital assistant a question, such as “Where am I right now?” Based on the user’s current location, the digital assistant answers, “You are in Central Park near the west gate.” The user also requests the performance of a task, for example, “Please invite my friends to my girlfriend’s birthday party next week.” In response, the digital assistant can acknowledge the request by saying “Yes, right away,” and then send a suitable calendar invite on behalf of the user to each of the user’s friends listed in the user’s electronic address book. During performance of a requested task, the digital assistant sometimes interacts with the user in a continuous dialogue involving multiple exchanges of information over an extended period of time. There are numerous other ways of interacting with a digital assistant to request information or performance of various

tasks. In addition to providing verbal responses and taking programmed actions, the digital assistant also provides responses in other visual or audio forms, e.g., as text, alerts, music, videos, animations, etc.

As shown in FIG. 1, in some examples, a digital assistant is implemented according to a client-server model. The digital assistant includes client-side portion 102 (hereafter “DA client 102”) executed on user device 104 and server-side portion 106 (hereafter “DA server 106”) executed on server system 108. DA client 102 communicates with DA server 106 through one or more networks 110. DA client 102 provides client-side functionalities such as user-facing input and output processing and communication with DA server 106. DA server 106 provides server-side functionalities for any number of DA clients 102 each residing on a respective user device 104.

In some examples, DA server 106 includes client-facing I/O interface 112, one or more processing modules 114, data and models 116, and I/O interface to external services 118. The client-facing I/O interface 112 facilitates the client-facing input and output processing for DA server 106. One or more processing modules 114 utilize data and models 116 to process speech input and determine the user’s intent based on natural language input. Further, one or more processing modules 114 perform task execution based on inferred user intent. In some examples, DA server 106 communicates with external services 120 through network(s) 110 for task completion or information acquisition. I/O interface to external services 118 facilitates such communications.

User device 104 can be any suitable electronic device. In some examples, user device 104 is a portable multifunctional device (e.g., device 200, described below with reference to FIG. 2A), a multifunctional device (e.g., device 400, described below with reference to FIG. 4), or a personal electronic device (e.g., device 600, described below with reference to FIGS. 6A-B). A portable multifunctional device is, for example, a mobile telephone that also contains other functions, such as PDA and/or music player functions. Specific examples of portable multifunction devices include the Apple Watch®, iPhone®, iPod Touch®, and iPad® devices from Apple Inc. of Cupertino, Calif. Other examples of portable multifunction devices include, without limitation, earphones/headphones, speakers, and laptop or tablet computers. Further, in some examples, user device 104 is a non-portable multifunctional device. In particular, user device 104 is a desktop computer, a game console, a speaker, a television, or a television set-top box. In some examples, user device 104 includes a touch-sensitive surface (e.g., touch screen displays and/or touchpads). Further, user device 104 optionally includes one or more other physical user-interface devices, such as a physical keyboard, a mouse, and/or a joystick. Various examples of electronic devices, such as multifunctional devices, are described below in greater detail.

Examples of communication network(s) 110 include local area networks (LAN) and wide area networks (WAN), e.g., the Internet. Communication network(s) 110 is implemented using any known network protocol, including various wired or wireless protocols, such as, for example, Ethernet, Universal Serial Bus (USB), FIREWIRE, Global System for Mobile Communications (GSM), Enhanced Data GSM Environment (EDGE), code division multiple access (CDMA), time division multiple access (TDMA), Bluetooth, Wi-Fi, voice over Internet Protocol (VoIP), Wi-MAX, or any other suitable communication protocol.

Server system 108 is implemented on one or more stand-alone data processing apparatus or a distributed network of

computers. In some examples, server system **108** also employs various virtual devices and/or services of third-party service providers (e.g., third-party cloud service providers) to provide the underlying computing resources and/or infrastructure resources of server system **108**.

In some examples, user device **104** communicates with DA server **106** via second user device **122**. Second user device **122** is similar or identical to user device **104**. For example, second user device **122** is similar to devices **200**, **400**, or **600** described below with reference to FIGS. **2A**, **4**, and **6A-B**. User device **104** is configured to communicatively couple to second user device **122** via a direct communication connection, such as Bluetooth, NFC, BTLE, or the like, or via a wired or wireless network, such as a local Wi-Fi network. In some examples, second user device **122** is configured to act as a proxy between user device **104** and DA server **106**. For example, DA client **102** of user device **104** is configured to transmit information (e.g., a user request received at user device **104**) to DA server **106** via second user device **122**. DA server **106** processes the information and returns relevant data (e.g., data content responsive to the user request) to user device **104** via second user device **122**.

In some examples, user device **104** is configured to communicate abbreviated requests for data to second user device **122** to reduce the amount of information transmitted from user device **104**. Second user device **122** is configured to determine supplemental information to add to the abbreviated request to generate a complete request to transmit to DA server **106**. This system architecture can advantageously allow user device **104** having limited communication capabilities and/or limited battery power (e.g., a watch or a similar compact electronic device) to access services provided by DA server **106** by using second user device **122**, having greater communication capabilities and/or battery power (e.g., a mobile phone, laptop computer, tablet computer, or the like), as a proxy to DA server **106**. While only two user devices **104** and **122** are shown in FIG. **1**, it should be appreciated that system **100**, in some examples, includes any number and type of user devices configured in this proxy configuration to communicate with DA server system **106**.

Although the digital assistant shown in FIG. **1** includes both a client-side portion (e.g., DA client **102**) and a server-side portion (e.g., DA server **106**), in some examples, the functions of a digital assistant are implemented as a standalone application installed on a user device. In addition, the divisions of functionalities between the client and server portions of the digital assistant can vary in different implementations. For instance, in some examples, the DA client is a thin-client that provides only user-facing input and output processing functions, and delegates all other functionalities of the digital assistant to a backend server.

2. Electronic Devices

Attention is now directed toward embodiments of electronic devices for implementing the client-side portion of a digital assistant. FIG. **2A** is a block diagram illustrating portable multifunction device **200** with touch-sensitive display system **212** in accordance with some embodiments. Touch-sensitive display **212** is sometimes called a “touch screen” for convenience and is sometimes known as or called a “touch-sensitive display system.” Device **200** includes memory **202** (which optionally includes one or more computer-readable storage mediums), memory controller **222**, one or more processing units (CPUs) **220**, peripherals interface **218**, RF circuitry **208**, audio circuitry **210**, speaker **211**, microphone **213**, input/output (I/O) subsystem **206**, other input control devices **216**, and external

port **224**. Device **200** optionally includes one or more optical sensors **264**. Device **200** optionally includes one or more contact intensity sensors **265** for detecting intensity of contacts on device **200** (e.g., a touch-sensitive surface such as touch-sensitive display system **212** of device **200**). Device **200** optionally includes one or more tactile output generators **267** for generating tactile outputs on device **200** (e.g., generating tactile outputs on a touch-sensitive surface such as touch-sensitive display system **212** of device **200** or touchpad **455** of device **400**). These components optionally communicate over one or more communication buses or signal lines **203**.

As used in the specification and claims, the term “intensity” of a contact on a touch-sensitive surface refers to the force or pressure (force per unit area) of a contact (e.g., a finger contact) on the touch-sensitive surface, or to a substitute (proxy) for the force or pressure of a contact on the touch-sensitive surface. The intensity of a contact has a range of values that includes at least four distinct values and more typically includes hundreds of distinct values (e.g., at least 256). Intensity of a contact is, optionally, determined (or measured) using various approaches and various sensors or combinations of sensors. For example, one or more force sensors underneath or adjacent to the touch-sensitive surface are, optionally, used to measure force at various points on the touch-sensitive surface. In some implementations, force measurements from multiple force sensors are combined (e.g., a weighted average) to determine an estimated force of a contact. Similarly, a pressure-sensitive tip of a stylus is, optionally, used to determine a pressure of the stylus on the touch-sensitive surface. Alternatively, the size of the contact area detected on the touch-sensitive surface and/or changes thereto, the capacitance of the touch-sensitive surface proximate to the contact and/or changes thereto, and/or the resistance of the touch-sensitive surface proximate to the contact and/or changes thereto are, optionally, used as a substitute for the force or pressure of the contact on the touch-sensitive surface. In some implementations, the substitute measurements for contact force or pressure are used directly to determine whether an intensity threshold has been exceeded (e.g., the intensity threshold is described in units corresponding to the substitute measurements). In some implementations, the substitute measurements for contact force or pressure are converted to an estimated force or pressure, and the estimated force or pressure is used to determine whether an intensity threshold has been exceeded (e.g., the intensity threshold is a pressure threshold measured in units of pressure). Using the intensity of a contact as an attribute of a user input allows for user access to additional device functionality that may otherwise not be accessible by the user on a reduced-size device with limited real estate for displaying affordances (e.g., on a touch-sensitive display) and/or receiving user input (e.g., via a touch-sensitive display, a touch-sensitive surface, or a physical/mechanical control such as a knob or a button).

As used in the specification and claims, the term “tactile output” refers to physical displacement of a device relative to a previous position of the device, physical displacement of a component (e.g., a touch-sensitive surface) of a device relative to another component (e.g., housing) of the device, or displacement of the component relative to a center of mass of the device that will be detected by a user with the user’s sense of touch. For example, in situations where the device or the component of the device is in contact with a surface of a user that is sensitive to touch (e.g., a finger, palm, or other part of a user’s hand), the tactile output generated by the physical displacement will be interpreted

by the user as a tactile sensation corresponding to a perceived change in physical characteristics of the device or the component of the device. For example, movement of a touch-sensitive surface (e.g., a touch-sensitive display or trackpad) is, optionally, interpreted by the user as a “down click” or “up click” of a physical actuator button. In some cases, a user will feel a tactile sensation such as an “down click” or “up click” even when there is no movement of a physical actuator button associated with the touch-sensitive surface that is physically pressed (e.g., displaced) by the user’s movements. As another example, movement of the touch-sensitive surface is, optionally, interpreted or sensed by the user as “roughness” of the touch-sensitive surface, even when there is no change in smoothness of the touch-sensitive surface. While such interpretations of touch by a user will be subject to the individualized sensory perceptions of the user, there are many sensory perceptions of touch that are common to a large majority of users. Thus, when a tactile output is described as corresponding to a particular sensory perception of a user (e.g., an “up click,” a “down click,” “roughness”), unless otherwise stated, the generated tactile output corresponds to physical displacement of the device or a component thereof that will generate the described sensory perception for a typical (or average) user.

It should be appreciated that device **200** is only one example of a portable multifunction device, and that device **200** optionally has more or fewer components than shown, optionally combines two or more components, or optionally has a different configuration or arrangement of the components. The various components shown in FIG. **2A** are implemented in hardware, software, or a combination of both hardware and software, including one or more signal processing and/or application-specific integrated circuits.

Memory **202** includes one or more computer-readable storage mediums. The computer-readable storage mediums are, for example, tangible and non-transitory. Memory **202** includes high-speed random access memory and also includes non-volatile memory, such as one or more magnetic disk storage devices, flash memory devices, or other non-volatile solid-state memory devices. Memory controller **222** controls access to memory **202** by other components of device **200**.

In some examples, a non-transitory computer-readable storage medium of memory **202** is used to store instructions (e.g., for performing aspects of processes described below) for use by or in connection with an instruction execution system, apparatus, or device, such as a computer-based system, processor-containing system, or other system that can fetch the instructions from the instruction execution system, apparatus, or device and execute the instructions. In other examples, the instructions (e.g., for performing aspects of the processes described below) are stored on a non-transitory computer-readable storage medium (not shown) of the server system **108** or are divided between the non-transitory computer-readable storage medium of memory **202** and the non-transitory computer-readable storage medium of server system **108**.

Peripherals interface **218** is used to couple input and output peripherals of the device to CPU **220** and memory **202**. The one or more processors **220** run or execute various software programs and/or sets of instructions stored in memory **202** to perform various functions for device **200** and to process data. In some embodiments, peripherals interface **218**, CPU **220**, and memory controller **222** are implemented on a single chip, such as chip **204**. In some other embodiments, they are implemented on separate chips.

RF (radio frequency) circuitry **208** receives and sends RF signals, also called electromagnetic signals. RF circuitry **208** converts electrical signals to/from electromagnetic signals and communicates with communications networks and other communications devices via the electromagnetic signals. RF circuitry **208** optionally includes well-known circuitry for performing these functions, including but not limited to an antenna system, an RF transceiver, one or more amplifiers, a tuner, one or more oscillators, a digital signal processor, a CODEC chipset, a subscriber identity module (SIM) card, memory, and so forth. RF circuitry **208** optionally communicates with networks, such as the Internet, also referred to as the World Wide Web (WWW), an intranet and/or a wireless network, such as a cellular telephone network, a wireless local area network (LAN) and/or a metropolitan area network (MAN), and other devices by wireless communication. The RF circuitry **208** optionally includes well-known circuitry for detecting near field communication (NFC) fields, such as by a short-range communication radio. The wireless communication optionally uses any of a plurality of communications standards, protocols, and technologies, including but not limited to Global System for Mobile Communications (GSM), Enhanced Data GSM Environment (EDGE), high-speed downlink packet access (HSDPA), high-speed uplink packet access (HSUPA), Evolution, Data-Only (EV-DO), HSPA, HSPA+, Dual-Cell HSPA (DC-HSPA), long term evolution (LTE), near field communication (NFC), wideband code division multiple access (W-CDMA), code division multiple access (CDMA), time division multiple access (TDMA), Bluetooth, Bluetooth Low Energy (BTLE), Wireless Fidelity (Wi-Fi) (e.g., IEEE 802.11a, IEEE 802.11b, IEEE 802.11g, IEEE 802.11n, and/or IEEE 802.11ac), voice over Internet Protocol (VoIP), Wi-MAX, a protocol for e mail (e.g., Internet message access protocol (IMAP) and/or post office protocol (POP)), instant messaging (e.g., extensible messaging and presence protocol (XMPP), Session Initiation Protocol for Instant Messaging and Presence Leveraging Extensions (SIMPLE), Instant Messaging and Presence Service (IMPS)), and/or Short Message Service (SMS), or any other suitable communication protocol, including communication protocols not yet developed as of the filing date of this document.

Audio circuitry **210**, speaker **211**, and microphone **213** provide an audio interface between a user and device **200**. Audio circuitry **210** receives audio data from peripherals interface **218**, converts the audio data to an electrical signal, and transmits the electrical signal to speaker **211**. Speaker **211** converts the electrical signal to human-audible sound waves. Audio circuitry **210** also receives electrical signals converted by microphone **213** from sound waves. Audio circuitry **210** converts the electrical signal to audio data and transmits the audio data to peripherals interface **218** for processing. Audio data are retrieved from and/or transmitted to memory **202** and/or RF circuitry **208** by peripherals interface **218**. In some embodiments, audio circuitry **210** also includes a headset jack (e.g., **312**, FIG. **3**). The headset jack provides an interface between audio circuitry **210** and removable audio input/output peripherals, such as output-only headphones or a headset with both output (e.g., a headphone for one or both ears) and input (e.g., a microphone).

I/O subsystem **206** couples input/output peripherals on device **200**, such as touch screen **212** and other input control devices **216**, to peripherals interface **218**. I/O subsystem **206** optionally includes display controller **256**, optical sensor controller **258**, intensity sensor controller **259**, haptic feedback controller **261**, and one or more input controllers **260**

for other input or control devices. The one or more input controllers **260** receive/send electrical signals from/to other input control devices **216**. The other input control devices **216** optionally include physical buttons (e.g., push buttons, rocker buttons, etc.), dials, slider switches, joysticks, click wheels, and so forth. In some alternate embodiments, input controller(s) **260** are, optionally, coupled to any (or none) of the following: a keyboard, an infrared port, a USB port, and a pointer device such as a mouse. The one or more buttons (e.g., **308**, FIG. 3) optionally include an up/down button for volume control of speaker **211** and/or microphone **213**. The one or more buttons optionally include a push button (e.g., **306**, FIG. 3).

A quick press of the push button disengages a lock of touch screen **212** or begin a process that uses gestures on the touch screen to unlock the device, as described in U.S. patent application Ser. No. 11/322,549, "Unlocking a Device by Performing Gestures on an Unlock Image," filed Dec. 23, 2005, U.S. Pat. No. 7,657,849, which is hereby incorporated by reference in its entirety. A longer press of the push button (e.g., **306**) turns power to device **200** on or off. The user is able to customize a functionality of one or more of the buttons. Touch screen **212** is used to implement virtual or soft buttons and one or more soft keyboards.

Touch-sensitive display **212** provides an input interface and an output interface between the device and a user. Display controller **256** receives and/or sends electrical signals from/to touch screen **212**. Touch screen **212** displays visual output to the user. The visual output includes graphics, text, icons, video, and any combination thereof (collectively termed "graphics"). In some embodiments, some or all of the visual output correspond to user-interface objects.

Touch screen **212** has a touch-sensitive surface, sensor, or set of sensors that accepts input from the user based on haptic and/or tactile contact. Touch screen **212** and display controller **256** (along with any associated modules and/or sets of instructions in memory **202**) detect contact (and any movement or breaking of the contact) on touch screen **212** and convert the detected contact into interaction with user-interface objects (e.g., one or more soft keys, icons, web pages, or images) that are displayed on touch screen **212**. In an exemplary embodiment, a point of contact between touch screen **212** and the user corresponds to a finger of the user.

Touch screen **212** uses LCD (liquid crystal display) technology, LPD (light emitting polymer display) technology, or LED (light emitting diode) technology, although other display technologies may be used in other embodiments. Touch screen **212** and display controller **256** detect contact and any movement or breaking thereof using any of a plurality of touch sensing technologies now known or later developed, including but not limited to capacitive, resistive, infrared, and surface acoustic wave technologies, as well as other proximity sensor arrays or other elements for determining one or more points of contact with touch screen **212**. In an exemplary embodiment, projected mutual capacitance sensing technology is used, such as that found in the iPhone® and iPod Touch® from Apple Inc. of Cupertino, Calif.

A touch-sensitive display in some embodiments of touch screen **212** is analogous to the multi-touch sensitive touchpads described in the following U.S. Pat. No. 6,323,846 (Westerman et al.), U.S. Pat. No. 6,570,557 (Westerman et al.), and/or U.S. Pat. No. 6,677,932 (Westerman), and/or U.S. Patent Publication 2002/0015024A1, each of which is hereby incorporated by reference in its entirety. However,

touch screen **212** displays visual output from device **200**, whereas touch-sensitive touchpads do not provide visual output.

A touch-sensitive display in some embodiments of touch screen **212** is as described in the following applications: (1) U.S. patent application Ser. No. 11/381,313, "Multipoint Touch Surface Controller," filed May 2, 2006; (2) U.S. patent application Ser. No. 10/840,862, "Multipoint Touchscreen," filed May 6, 2004; (3) U.S. patent application Ser. No. 10/903,964, "Gestures For Touch Sensitive Input Devices," filed Jul. 30, 2004; (4) U.S. patent application Ser. No. 11/048,264, "Gestures For Touch Sensitive Input Devices," filed Jan. 31, 2005; (5) U.S. patent application Ser. No. 11/038,590, "Mode-Based Graphical User Interfaces For Touch Sensitive Input Devices," filed Jan. 18, 2005; (6) U.S. patent application Ser. No. 11/228,758, "Virtual Input Device Placement On A Touch Screen User Interface," filed Sep. 16, 2005; (7) U.S. patent application Ser. No. 11/228,700, "Operation Of A Computer With A Touch Screen Interface," filed Sep. 16, 2005; (8) U.S. patent application Ser. No. 11/228,737, "Activating Virtual Keys Of A Touch-Screen Virtual Keyboard," filed Sep. 16, 2005; and (9) U.S. patent application Ser. No. 11/367,749, "Multi-Functional Hand-Held Device," filed Mar. 3, 2006. All of these applications are incorporated by reference herein in their entirety.

Touch screen **212** has, for example, a video resolution in excess of 100 dpi. In some embodiments, the touch screen has a video resolution of approximately 160 dpi. The user makes contact with touch screen **212** using any suitable object or appendage, such as a stylus, a finger, and so forth. In some embodiments, the user interface is designed to work primarily with finger-based contacts and gestures, which can be less precise than stylus-based input due to the larger area of contact of a finger on the touch screen. In some embodiments, the device translates the rough finger-based input into a precise pointer/cursor position or command for performing the actions desired by the user.

In some embodiments, in addition to the touch screen, device **200** includes a touchpad (not shown) for activating or deactivating particular functions. In some embodiments, the touchpad is a touch-sensitive area of the device that, unlike the touch screen, does not display visual output. The touchpad is a touch-sensitive surface that is separate from touch screen **212** or an extension of the touch-sensitive surface formed by the touch screen.

Device **200** also includes power system **262** for powering the various components. Power system **262** includes a power management system, one or more power sources (e.g., battery, alternating current (AC)), a recharging system, a power failure detection circuit, a power converter or inverter, a power status indicator (e.g., a light-emitting diode (LED)) and any other components associated with the generation, management and distribution of power in portable devices.

Device **200** also includes one or more optical sensors **264**. FIG. 2A shows an optical sensor coupled to optical sensor controller **258** in I/O subsystem **206**. Optical sensor **264** includes charge-coupled device (CCD) or complementary metal-oxide semiconductor (CMOS) phototransistors. Optical sensor **264** receives light from the environment, projected through one or more lenses, and converts the light to data representing an image. In conjunction with imaging module **243** (also called a camera module), optical sensor **264** captures still images or video. In some embodiments, an optical sensor is located on the back of device **200**, opposite touch screen display **212** on the front of the device so that the touch screen display is used as a viewfinder for still

and/or video image acquisition. In some embodiments, an optical sensor is located on the front of the device so that the user's image is obtained for video conferencing while the user views the other video conference participants on the touch screen display. In some embodiments, the position of optical sensor 264 can be changed by the user (e.g., by rotating the lens and the sensor in the device housing) so that a single optical sensor 264 is used along with the touch screen display for both video conferencing and still and/or video image acquisition.

Device 200 optionally also includes one or more contact intensity sensors 265. FIG. 2A shows a contact intensity sensor coupled to intensity sensor controller 259 in I/O subsystem 206. Contact intensity sensor 265 optionally includes one or more piezoresistive strain gauges, capacitive force sensors, electric force sensors, piezoelectric force sensors, optical force sensors, capacitive touch-sensitive surfaces, or other intensity sensors (e.g., sensors used to measure the force (or pressure) of a contact on a touch-sensitive surface). Contact intensity sensor 265 receives contact intensity information (e.g., pressure information or a proxy for pressure information) from the environment. In some embodiments, at least one contact intensity sensor is collocated with, or proximate to, a touch-sensitive surface (e.g., touch-sensitive display system 212). In some embodiments, at least one contact intensity sensor is located on the back of device 200, opposite touch screen display 212, which is located on the front of device 200.

Device 200 also includes one or more proximity sensors 266. FIG. 2A shows proximity sensor 266 coupled to peripherals interface 218. Alternately, proximity sensor 266 is coupled to input controller 260 in I/O subsystem 206. Proximity sensor 266 is performed as described in U.S. patent application Ser. No. 11/241,839, "Proximity Detector In Handheld Device"; Ser. No. 11/240,788, "Proximity Detector In Handheld Device"; Ser. No. 11/620,702, "Using Ambient Light Sensor To Augment Proximity Sensor Output"; Ser. No. 11/586,862, "Automated Response To And Sensing Of User Activity In Portable Devices"; and Ser. No. 11/638,251, "Methods And Systems For Automatic Configuration Of Peripherals," which are hereby incorporated by reference in their entirety. In some embodiments, the proximity sensor turns off and disables touch screen 212 when the multifunction device is placed near the user's ear (e.g., when the user is making a phone call).

Device 200 optionally also includes one or more tactile output generators 267. FIG. 2A shows a tactile output generator coupled to haptic feedback controller 261 in I/O subsystem 206. Tactile output generator 267 optionally includes one or more electroacoustic devices such as speakers or other audio components and/or electromechanical devices that convert energy into linear motion such as a motor, solenoid, electroactive polymer, piezoelectric actuator, electrostatic actuator, or other tactile output generating component (e.g., a component that converts electrical signals into tactile outputs on the device). Contact intensity sensor 265 receives tactile feedback generation instructions from haptic feedback module 233 and generates tactile outputs on device 200 that are capable of being sensed by a user of device 200. In some embodiments, at least one tactile output generator is collocated with, or proximate to, a touch-sensitive surface (e.g., touch-sensitive display system 212) and, optionally, generates a tactile output by moving the touch-sensitive surface vertically (e.g., in/out of a surface of device 200) or laterally (e.g., back and forth in the same plane as a surface of device 200). In some embodiments, at least one tactile output generator sensor is located

on the back of device 200, opposite touch screen display 212, which is located on the front of device 200.

Device 200 also includes one or more accelerometers 268. FIG. 2A shows accelerometer 268 coupled to peripherals interface 218. Alternately, accelerometer 268 is coupled to an input controller 260 in I/O subsystem 206. Accelerometer 268 performs, for example, as described in U.S. Patent Publication No. 20050190059, "Acceleration-based Theft Detection System for Portable Electronic Devices," and U.S. Patent Publication No. 20060017692, "Methods And Apparatuses For Operating A Portable Device Based On An Accelerometer," both of which are incorporated by reference herein in their entirety. In some embodiments, information is displayed on the touch screen display in a portrait view or a landscape view based on an analysis of data received from the one or more accelerometers. Device 200 optionally includes, in addition to accelerometer(s) 268, a magnetometer (not shown) and a GPS (or GLONASS or other global navigation system) receiver (not shown) for obtaining information concerning the location and orientation (e.g., portrait or landscape) of device 200.

In some embodiments, the software components stored in memory 202 include operating system 226, communication module (or set of instructions) 228, contact/motion module (or set of instructions) 230, graphics module (or set of instructions) 232, text input module (or set of instructions) 234, Global Positioning System (GPS) module (or set of instructions) 235, Digital Assistant Client Module 229, and applications (or sets of instructions) 236. Further, memory 202 stores data and models, such as user data and models 231. Furthermore, in some embodiments, memory 202 (FIG. 2A) or 470 (FIG. 4) stores device/global internal state 257, as shown in FIGS. 2A and 4. Device/global internal state 257 includes one or more of: active application state, indicating which applications, if any, are currently active; display state, indicating what applications, views or other information occupy various regions of touch screen display 212; sensor state, including information obtained from the device's various sensors and input control devices 216; and location information concerning the device's location and/or attitude.

Operating system 226 (e.g., Darwin, RTXC, LINUX, UNIX, OS X, iOS, WINDOWS, or an embedded operating system such as VxWorks) includes various software components and/or drivers for controlling and managing general system tasks (e.g., memory management, storage device control, power management, etc.) and facilitates communication between various hardware and software components.

Communication module 228 facilitates communication with other devices over one or more external ports 224 and also includes various software components for handling data received by RF circuitry 208 and/or external port 224. External port 224 (e.g., Universal Serial Bus (USB), FIREWIRE, etc.) is adapted for coupling directly to other devices or indirectly over a network (e.g., the Internet, wireless LAN, etc.). In some embodiments, the external port is a multi-pin (e.g., 30-pin) connector that is the same as, or similar to and/or compatible with, the 30-pin connector used on iPod® (trademark of Apple Inc.) devices.

Contact/motion module 230 optionally detects contact with touch screen 212 (in conjunction with display controller 256) and other touch-sensitive devices (e.g., a touchpad or physical click wheel). Contact/motion module 230 includes various software components for performing various operations related to detection of contact, such as determining if contact has occurred (e.g., detecting a finger-down event), determining an intensity of the contact (e.g., the force or pressure of the contact or a substitute for the

force or pressure of the contact), determining if there is movement of the contact and tracking the movement across the touch-sensitive surface (e.g., detecting one or more finger-dragging events), and determining if the contact has ceased (e.g., detecting a finger-up event or a break in contact). Contact/motion module **230** receives contact data from the touch-sensitive surface. Determining movement of the point of contact, which is represented by a series of contact data, optionally includes determining speed (magnitude), velocity (magnitude and direction), and/or an acceleration (a change in magnitude and/or direction) of the point of contact. These operations are, optionally, applied to single contacts (e.g., one finger contacts) or to multiple simultaneous contacts (e.g., “multitouch”/multiple finger contacts). In some embodiments, contact/motion module **230** and display controller **256** detect contact on a touchpad.

In some embodiments, contact/motion module **230** uses a set of one or more intensity thresholds to determine whether an operation has been performed by a user (e.g., to determine whether a user has “clicked” on an icon). In some embodiments, at least a subset of the intensity thresholds are determined in accordance with software parameters (e.g., the intensity thresholds are not determined by the activation thresholds of particular physical actuators and can be adjusted without changing the physical hardware of device **200**). For example, a mouse “click” threshold of a trackpad or touch screen display can be set to any of a large range of predefined threshold values without changing the trackpad or touch screen display hardware. Additionally, in some implementations, a user of the device is provided with software settings for adjusting one or more of the set of intensity thresholds (e.g., by adjusting individual intensity thresholds and/or by adjusting a plurality of intensity thresholds at once with a system-level click “intensity” parameter).

Contact/motion module **230** optionally detects a gesture input by a user. Different gestures on the touch-sensitive surface have different contact patterns (e.g., different motions, timings, and/or intensities of detected contacts). Thus, a gesture is, optionally, detected by detecting a particular contact pattern. For example, detecting a finger tap gesture includes detecting a finger-down event followed by detecting a finger-up (liftoff) event at the same position (or substantially the same position) as the finger-down event (e.g., at the position of an icon). As another example, detecting a finger swipe gesture on the touch-sensitive surface includes detecting a finger-down event followed by detecting one or more finger-dragging events, and subsequently followed by detecting a finger-up (liftoff) event.

Graphics module **232** includes various known software components for rendering and displaying graphics on touch screen **212** or other display, including components for changing the visual impact (e.g., brightness, transparency, saturation, contrast, or other visual property) of graphics that are displayed. As used herein, the term “graphics” includes any object that can be displayed to a user, including, without limitation, text, web pages, icons (such as user-interface objects including soft keys), digital images, videos, animations, and the like.

In some embodiments, graphics module **232** stores data representing graphics to be used. Each graphic is, optionally, assigned a corresponding code. Graphics module **232** receives, from applications etc., one or more codes specifying graphics to be displayed along with, if necessary, coordinate data and other graphic property data, and then generates screen image data to output to display controller **256**.

Haptic feedback module **233** includes various software components for generating instructions used by tactile output generator(s) **267** to produce tactile outputs at one or more locations on device **200** in response to user interactions with device **200**.

Text input module **234**, which is, in some examples, a component of graphics module **232**, provides soft keyboards for entering text in various applications (e.g., contacts **237**, email **240**, IM **241**, browser **247**, and any other application that needs text input).

GPS module **235** determines the location of the device and provides this information for use in various applications (e.g., to telephone **238** for use in location-based dialing; to camera **243** as picture/video metadata; and to applications that provide location-based services such as weather widgets, local yellow page widgets, and map/navigation widgets).

Digital assistant client module **229** includes various client-side digital assistant instructions to provide the client-side functionalities of the digital assistant. For example, digital assistant client module **229** is capable of accepting voice input (e.g., speech input), text input, touch input, and/or gestural input through various user interfaces (e.g., microphone **213**, accelerometer(s) **268**, touch-sensitive display system **212**, optical sensor(s) **264**, other input control devices **216**, etc.) of portable multifunction device **200**. Digital assistant client module **229** is also capable of providing output in audio (e.g., speech output), visual, and/or tactile forms through various output interfaces (e.g., speaker **211**, touch-sensitive display system **212**, tactile output generator(s) **267**, etc.) of portable multifunction device **200**. For example, output is provided as voice, sound, alerts, text messages, menus, graphics, videos, animations, vibrations, and/or combinations of two or more of the above. During operation, digital assistant client module **229** communicates with DA server **106** using RF circuitry **208**.

User data and models **231** include various data associated with the user (e.g., user-specific vocabulary data, user preference data, user-specified name pronunciations, data from the user’s electronic address book, to-do lists, shopping lists, etc.) to provide the client-side functionalities of the digital assistant. Further, user data and models **231** include various models (e.g., speech recognition models, statistical language models, natural language processing models, ontology, task flow models, service models, etc.) for processing user input and determining user intent.

In some examples, digital assistant client module **229** utilizes the various sensors, subsystems, and peripheral devices of portable multifunction device **200** to gather additional information from the surrounding environment of the portable multifunction device **200** to establish a context associated with a user, the current user interaction, and/or the current user input. In some examples, digital assistant client module **229** provides the contextual information or a subset thereof with the user input to DA server **106** to help infer the user’s intent. In some examples, the digital assistant also uses the contextual information to determine how to prepare and deliver outputs to the user. Contextual information is referred to as context data.

In some examples, the contextual information that accompanies the user input includes sensor information, e.g., lighting, ambient noise, ambient temperature, images or videos of the surrounding environment, etc. In some examples, the contextual information can also include the physical state of the device, e.g., device orientation, device location, device temperature, power level, speed, acceleration, motion patterns, cellular signals strength, etc. In some

examples, information related to the software state of DA server 106, e.g., running processes, installed programs, past and present network activities, background services, error logs, resources usage, etc., and of portable multifunction device 200 is provided to DA server 106 as contextual information associated with a user input.

In some examples, the digital assistant client module 229 selectively provides information (e.g., user data 231) stored on the portable multifunction device 200 in response to requests from DA server 106. In some examples, digital assistant client module 229 also elicits additional input from the user via a natural language dialogue or other user interfaces upon request by DA server 106. Digital assistant client module 229 passes the additional input to DA server 106 to help DA server 106 in intent deduction and/or fulfillment of the user's intent expressed in the user request.

A more detailed description of a digital assistant is described below with reference to FIGS. 7A-C. It should be recognized that digital assistant client module 229 can include any number of the sub-modules of digital assistant module 726 described below.

Applications 236 include the following modules (or sets of instructions), or a subset or superset thereof:

Contacts module 237 (sometimes called an address book or contact list);

Telephone module 238;

Video conference module 239;

E-mail client module 240;

Instant messaging (IM) module 241;

Workout support module 242;

Camera module 243 for still and/or video images;

Image management module 244;

Video player module;

Music player module;

Browser module 247;

Calendar module 248;

Widget modules 249, which includes, in some examples, one or more of: weather widget 249-1, stocks widget 249-2, calculator widget 249-3, alarm clock widget 249-4, dictionary widget 249-5, and other widgets obtained by the user, as well as user-created widgets 249-6;

Widget creator module 250 for making user-created widgets 249-6;

Search module 251;

Video and music player module 252, which merges video player module and music player module;

Notes module 253;

Map module 254; and/or

Online video module 255.

Examples of other applications 236 that are stored in memory 202 include other word processing applications, other image editing applications, drawing applications, presentation applications, JAVA-enabled applications, encryption, digital rights management, voice recognition, and voice replication.

In conjunction with touch screen 212, display controller 256, contact/motion module 230, graphics module 232, and text input module 234, contacts module 237 are used to manage an address book or contact list (e.g., stored in application internal state 292 of contacts module 237 in memory 202 or memory 470), including: adding name(s) to the address book; deleting name(s) from the address book; associating telephone number(s), e-mail address(es), physical address(es) or other information with a name; associating an image with a name; categorizing and sorting names; providing telephone numbers or e-mail addresses to initiate

and/or facilitate communications by telephone 238, video conference module 239, e-mail 240, or IM 241; and so forth.

In conjunction with RF circuitry 208, audio circuitry 210, speaker 211, microphone 213, touch screen 212, display controller 256, contact/motion module 230, graphics module 232, and text input module 234, telephone module 238 are used to enter a sequence of characters corresponding to a telephone number, access one or more telephone numbers in contacts module 237, modify a telephone number that has been entered, dial a respective telephone number, conduct a conversation, and disconnect or hang up when the conversation is completed. As noted above, the wireless communication uses any of a plurality of communications standards, protocols, and technologies.

In conjunction with RF circuitry 208, audio circuitry 210, speaker 211, microphone 213, touch screen 212, display controller 256, optical sensor 264, optical sensor controller 258, contact/motion module 230, graphics module 232, text input module 234, contacts module 237, and telephone module 238, video conference module 239 includes executable instructions to initiate, conduct, and terminate a video conference between a user and one or more other participants in accordance with user instructions.

In conjunction with RF circuitry 208, touch screen 212, display controller 256, contact/motion module 230, graphics module 232, and text input module 234, e-mail client module 240 includes executable instructions to create, send, receive, and manage e-mail in response to user instructions. In conjunction with image management module 244, e-mail client module 240 makes it very easy to create and send e-mails with still or video images taken with camera module 243.

In conjunction with RF circuitry 208, touch screen 212, display controller 256, contact/motion module 230, graphics module 232, and text input module 234, the instant messaging module 241 includes executable instructions to enter a sequence of characters corresponding to an instant message, to modify previously entered characters, to transmit a respective instant message (for example, using a Short Message Service (SMS) or Multimedia Message Service (MMS) protocol for telephony-based instant messages or using XMPP, SIMPLE, or IMPS for Internet-based instant messages), to receive instant messages, and to view received instant messages. In some embodiments, transmitted and/or received instant messages include graphics, photos, audio files, video files and/or other attachments as are supported in an MMS and/or an Enhanced Messaging Service (EMS). As used herein, "instant messaging" refers to both telephony-based messages (e.g., messages sent using SMS or MMS) and Internet-based messages (e.g., messages sent using XMPP, SIMPLE, or IMPS).

In conjunction with RF circuitry 208, touch screen 212, display controller 256, contact/motion module 230, graphics module 232, text input module 234, GPS module 235, map module 254, and music player module, workout support module 242 includes executable instructions to create workouts (e.g., with time, distance, and/or calorie burning goals); communicate with workout sensors (sports devices); receive workout sensor data; calibrate sensors used to monitor a workout; select and play music for a workout; and display, store, and transmit workout data.

In conjunction with touch screen 212, display controller 256, optical sensor(s) 264, optical sensor controller 258, contact/motion module 230, graphics module 232, and image management module 244, camera module 243 includes executable instructions to capture still images or video (including a video stream) and store them into

memory 202, modify characteristics of a still image or video, or delete a still image or video from memory 202.

In conjunction with touch screen 212, display controller 256, contact/motion module 230, graphics module 232, text input module 234, and camera module 243, image management module 244 includes executable instructions to arrange, modify (e.g., edit), or otherwise manipulate, label, delete, present (e.g., in a digital slide show or album), and store still and/or video images.

In conjunction with RF circuitry 208, touch screen 212, display controller 256, contact/motion module 230, graphics module 232, and text input module 234, browser module 247 includes executable instructions to browse the Internet in accordance with user instructions, including searching, linking to, receiving, and displaying web pages or portions thereof, as well as attachments and other files linked to web pages.

In conjunction with RF circuitry 208, touch screen 212, display controller 256, contact/motion module 230, graphics module 232, text input module 234, e-mail client module 240, and browser module 247, calendar module 248 includes executable instructions to create, display, modify, and store calendars and data associated with calendars (e.g., calendar entries, to-do lists, etc.) in accordance with user instructions.

In conjunction with RF circuitry 208, touch screen 212, display controller 256, contact/motion module 230, graphics module 232, text input module 234, and browser module 247, widget modules 249 are mini-applications that can be downloaded and used by a user (e.g., weather widget 249-1, stocks widget 249-2, calculator widget 249-3, alarm clock widget 249-4, and dictionary widget 249-5) or created by the user (e.g., user-created widget 249-6). In some embodiments, a widget includes an HTML (Hypertext Markup Language) file, a CSS (Cascading Style Sheets) file, and a JavaScript file. In some embodiments, a widget includes an XML (Extensible Markup Language) file and a JavaScript file (e.g., Yahoo! Widgets).

In conjunction with RF circuitry 208, touch screen 212, display controller 256, contact/motion module 230, graphics module 232, text input module 234, and browser module 247, the widget creator module 250 are used by a user to create widgets (e.g., turning a user-specified portion of a web page into a widget).

In conjunction with touch screen 212, display controller 256, contact/motion module 230, graphics module 232, and text input module 234, search module 251 includes executable instructions to search for text, music, sound, image, video, and/or other files in memory 202 that match one or more search criteria (e.g., one or more user-specified search terms) in accordance with user instructions.

In conjunction with touch screen 212, display controller 256, contact/motion module 230, graphics module 232, audio circuitry 210, speaker 211, RF circuitry 208, and browser module 247, video and music player module 252 includes executable instructions that allow the user to download and play back recorded music and other sound files stored in one or more file formats, such as MP3 or AAC files, and executable instructions to display, present, or otherwise play back videos (e.g., on touch screen 212 or on an external, connected display via external port 224). In some embodiments, device 200 optionally includes the functionality of an MP3 player, such as an iPod (trademark of Apple Inc.).

In conjunction with touch screen 212, display controller 256, contact/motion module 230, graphics module 232, and text input module 234, notes module 253 includes executable instructions to create and manage notes, to-do lists, and the like in accordance with user instructions.

In conjunction with RF circuitry 208, touch screen 212, display controller 256, contact/motion module 230, graphics module 232, text input module 234, GPS module 235, and browser module 247, map module 254 are used to receive, display, modify, and store maps and data associated with maps (e.g., driving directions, data on stores and other points of interest at or near a particular location, and other location-based data) in accordance with user instructions.

In conjunction with touch screen 212, display controller 256, contact/motion module 230, graphics module 232, audio circuitry 210, speaker 211, RF circuitry 208, text input module 234, e-mail client module 240, and browser module 247, online video module 255 includes instructions that allow the user to access, browse, receive (e.g., by streaming and/or download), play back (e.g., on the touch screen or on an external, connected display via external port 224), send an e-mail with a link to a particular online video, and otherwise manage online videos in one or more file formats, such as H.264. In some embodiments, instant messaging module 241, rather than e-mail client module 240, is used to send a link to a particular online video. Additional description of the online video application can be found in U.S. Provisional Patent Application No. 60/936,562, "Portable Multifunction Device, Method, and Graphical User Interface for Playing Online Videos," filed Jun. 20, 2007, and U.S. patent application Ser. No. 11/968,067, "Portable Multifunction Device, Method, and Graphical User Interface for Playing Online Videos," filed Dec. 31, 2007, the contents of which are hereby incorporated by reference in their entirety.

Each of the above-identified modules and applications corresponds to a set of executable instructions for performing one or more functions described above and the methods described in this application (e.g., the computer-implemented methods and other information processing methods described herein). These modules (e.g., sets of instructions) need not be implemented as separate software programs, procedures, or modules, and thus various subsets of these modules can be combined or otherwise rearranged in various embodiments. For example, video player module can be combined with music player module into a single module (e.g., video and music player module 252, FIG. 2A). In some embodiments, memory 202 stores a subset of the modules and data structures identified above. Furthermore, memory 202 stores additional modules and data structures not described above.

In some embodiments, device 200 is a device where operation of a predefined set of functions on the device is performed exclusively through a touch screen and/or a touchpad. By using a touch screen and/or a touchpad as the primary input control device for operation of device 200, the number of physical input control devices (such as push buttons, dials, and the like) on device 200 is reduced.

The predefined set of functions that are performed exclusively through a touch screen and/or a touchpad optionally include navigation between user interfaces. In some embodiments, the touchpad, when touched by the user, navigates device 200 to a main, home, or root menu from any user interface that is displayed on device 200. In such embodiments, a "menu button" is implemented using a touchpad. In some other embodiments, the menu button is a physical push button or other physical input control device instead of a touchpad.

FIG. 2B is a block diagram illustrating exemplary components for event handling in accordance with some embodiments. In some embodiments, memory 202 (FIG. 2A) or 470 (FIG. 4) includes event sorter 270 (e.g., in

operating system 226) and a respective application 236-1 (e.g., any of the aforementioned applications 237-251, 255, 480-490).

Event sorter 270 receives event information and determines the application 236-1 and application view 291 of application 236-1 to which to deliver the event information. Event sorter 270 includes event monitor 271 and event dispatcher module 274. In some embodiments, application 236-1 includes application internal state 292, which indicates the current application view(s) displayed on touch-sensitive display 212 when the application is active or executing. In some embodiments, device/global internal state 257 is used by event sorter 270 to determine which application(s) is (are) currently active, and application internal state 292 is used by event sorter 270 to determine application views 291 to which to deliver event information.

In some embodiments, application internal state 292 includes additional information, such as one or more of: resume information to be used when application 236-1 resumes execution, user interface state information that indicates information being displayed or that is ready for display by application 236-1, a state queue for enabling the user to go back to a prior state or view of application 236-1, and a redo/undo queue of previous actions taken by the user.

Event monitor 271 receives event information from peripherals interface 218. Event information includes information about a sub-event (e.g., a user touch on touch-sensitive display 212, as part of a multi-touch gesture). Peripherals interface 218 transmits information it receives from I/O subsystem 206 or a sensor, such as proximity sensor 266, accelerometer(s) 268, and/or microphone 213 (through audio circuitry 210). Information that peripherals interface 218 receives from I/O subsystem 206 includes information from touch-sensitive display 212 or a touch-sensitive surface.

In some embodiments, event monitor 271 sends requests to the peripherals interface 218 at predetermined intervals. In response, peripherals interface 218 transmits event information. In other embodiments, peripherals interface 218 transmits event information only when there is a significant event (e.g., receiving an input above a predetermined noise threshold and/or for more than a predetermined duration).

In some embodiments, event sorter 270 also includes a hit view determination module 272 and/or an active event recognizer determination module 273.

Hit view determination module 272 provides software procedures for determining where a sub-event has taken place within one or more views when touch-sensitive display 212 displays more than one view. Views are made up of controls and other elements that a user can see on the display.

Another aspect of the user interface associated with an application is a set of views, sometimes herein called application views or user interface windows, in which information is displayed and touch-based gestures occur. The application views (of a respective application) in which a touch is detected correspond to programmatic levels within a programmatic or view hierarchy of the application. For example, the lowest level view in which a touch is detected is called the hit view, and the set of events that are recognized as proper inputs is determined based, at least in part, on the hit view of the initial touch that begins a touch-based gesture.

Hit view determination module 272 receives information related to sub events of a touch-based gesture. When an application has multiple views organized in a hierarchy, hit view determination module 272 identifies a hit view as the

lowest view in the hierarchy which should handle the sub-event. In most circumstances, the hit view is the lowest level view in which an initiating sub-event occurs (e.g., the first sub-event in the sequence of sub-events that form an event or potential event). Once the hit view is identified by the hit view determination module 272, the hit view typically receives all sub-events related to the same touch or input source for which it was identified as the hit view.

Active event recognizer determination module 273 determines which view or views within a view hierarchy should receive a particular sequence of sub-events. In some embodiments, active event recognizer determination module 273 determines that only the hit view should receive a particular sequence of sub-events. In other embodiments, active event recognizer determination module 273 determines that all views that include the physical location of a sub-event are actively involved views, and therefore determines that all actively involved views should receive a particular sequence of sub-events. In other embodiments, even if touch sub-events were entirely confined to the area associated with one particular view, views higher in the hierarchy would still remain as actively involved views.

Event dispatcher module 274 dispatches the event information to an event recognizer (e.g., event recognizer 280). In embodiments including active event recognizer determination module 273, event dispatcher module 274 delivers the event information to an event recognizer determined by active event recognizer determination module 273. In some embodiments, event dispatcher module 274 stores in an event queue the event information, which is retrieved by a respective event receiver 282.

In some embodiments, operating system 226 includes event sorter 270. Alternatively, application 236-1 includes event sorter 270. In yet other embodiments, event sorter 270 is a stand-alone module, or a part of another module stored in memory 202, such as contact/motion module 230.

In some embodiments, application 236-1 includes a plurality of event handlers 290 and one or more application views 291, each of which includes instructions for handling touch events that occur within a respective view of the application's user interface. Each application view 291 of the application 236-1 includes one or more event recognizers 280. Typically, a respective application view 291 includes a plurality of event recognizers 280. In other embodiments, one or more of event recognizers 280 are part of a separate module, such as a user interface kit (not shown) or a higher level object from which application 236-1 inherits methods and other properties. In some embodiments, a respective event handler 290 includes one or more of: data updater 276, object updater 277, GUI updater 278, and/or event data 279 received from event sorter 270. Event handler 290 utilizes or calls data updater 276, object updater 277, or GUI updater 278 to update the application internal state 292. Alternatively, one or more of the application views 291 include one or more respective event handlers 290. Also, in some embodiments, one or more of data updater 276, object updater 277, and GUI updater 278 are included in a respective application view 291.

A respective event recognizer 280 receives event information (e.g., event data 279) from event sorter 270 and identifies an event from the event information. Event recognizer 280 includes event receiver 282 and event comparator 284. In some embodiments, event recognizer 280 also includes at least a subset of: metadata 283, and event delivery instructions 288 (which include sub-event delivery instructions).

Event receiver **282** receives event information from event sorter **270**. The event information includes information about a sub-event, for example, a touch or a touch movement. Depending on the sub-event, the event information also includes additional information, such as location of the sub-event. When the sub-event concerns motion of a touch, the event information also includes speed and direction of the sub-event. In some embodiments, events include rotation of the device from one orientation to another (e.g., from a portrait orientation to a landscape orientation, or vice versa), and the event information includes corresponding information about the current orientation (also called device attitude) of the device.

Event comparator **284** compares the event information to predefined event or sub-event definitions and, based on the comparison, determines an event or sub event, or determines or updates the state of an event or sub-event. In some embodiments, event comparator **284** includes event definitions **286**. Event definitions **286** contain definitions of events (e.g., predefined sequences of sub-events), for example, event **1** (**287-1**), event **2** (**287-2**), and others. In some embodiments, sub-events in an event (**287**) include, for example, touch begin, touch end, touch movement, touch cancellation, and multiple touching. In one example, the definition for event **1** (**287-1**) is a double tap on a displayed object. The double tap, for example, comprises a first touch (touch begin) on the displayed object for a predetermined phase, a first liftoff (touch end) for a predetermined phase, a second touch (touch begin) on the displayed object for a predetermined phase, and a second liftoff (touch end) for a predetermined phase. In another example, the definition for event **2** (**287-2**) is a dragging on a displayed object. The dragging, for example, comprises a touch (or contact) on the displayed object for a predetermined phase, a movement of the touch across touch-sensitive display **212**, and liftoff of the touch (touch end). In some embodiments, the event also includes information for one or more associated event handlers **290**.

In some embodiments, event definition **287** includes a definition of an event for a respective user-interface object. In some embodiments, event comparator **284** performs a hit test to determine which user-interface object is associated with a sub-event. For example, in an application view in which three user-interface objects are displayed on touch-sensitive display **212**, when a touch is detected on touch-sensitive display **212**, event comparator **284** performs a hit test to determine which of the three user-interface objects is associated with the touch (sub-event). If each displayed object is associated with a respective event handler **290**, the event comparator uses the result of the hit test to determine which event handler **290** should be activated. For example, event comparator **284** selects an event handler associated with the sub-event and the object triggering the hit test.

In some embodiments, the definition for a respective event (**287**) also includes delayed actions that delay delivery of the event information until after it has been determined whether the sequence of sub-events does or does not correspond to the event recognizer's event type.

When a respective event recognizer **280** determines that the series of sub-events do not match any of the events in event definitions **286**, the respective event recognizer **280** enters an event impossible, event failed, or event ended state, after which it disregards subsequent sub-events of the touch-based gesture. In this situation, other event recognizers, if any, that remain active for the hit view continue to track and process sub-events of an ongoing touch-based gesture.

In some embodiments, a respective event recognizer **280** includes metadata **283** with configurable properties, flags, and/or lists that indicate how the event delivery system should perform sub-event delivery to actively involved event recognizers. In some embodiments, metadata **283** includes configurable properties, flags, and/or lists that indicate how event recognizers interact, or are enabled to interact, with one another. In some embodiments, metadata **283** includes configurable properties, flags, and/or lists that indicate whether sub-events are delivered to varying levels in the view or programmatic hierarchy.

In some embodiments, a respective event recognizer **280** activates event handler **290** associated with an event when one or more particular sub-events of an event are recognized. In some embodiments, a respective event recognizer **280** delivers event information associated with the event to event handler **290**. Activating an event handler **290** is distinct from sending (and deferred sending) sub-events to a respective hit view. In some embodiments, event recognizer **280** throws a flag associated with the recognized event, and event handler **290** associated with the flag catches the flag and performs a predefined process.

In some embodiments, event delivery instructions **288** include sub-event delivery instructions that deliver event information about a sub-event without activating an event handler. Instead, the sub-event delivery instructions deliver event information to event handlers associated with the series of sub-events or to actively involved views. Event handlers associated with the series of sub-events or with actively involved views receive the event information and perform a predetermined process.

In some embodiments, data updater **276** creates and updates data used in application **236-1**. For example, data updater **276** updates the telephone number used in contacts module **237**, or stores a video file used in video player module. In some embodiments, object updater **277** creates and updates objects used in application **236-1**. For example, object updater **277** creates a new user-interface object or updates the position of a user-interface object. GUI updater **278** updates the GUI. For example, GUI updater **278** prepares display information and sends it to graphics module **232** for display on a touch-sensitive display.

In some embodiments, event handler(s) **290** includes or has access to data updater **276**, object updater **277**, and GUI updater **278**. In some embodiments, data updater **276**, object updater **277**, and GUI updater **278** are included in a single module of a respective application **236-1** or application view **291**. In other embodiments, they are included in two or more software modules.

It shall be understood that the foregoing discussion regarding event handling of user touches on touch-sensitive displays also applies to other forms of user inputs to operate multifunction devices **200** with input devices, not all of which are initiated on touch screens. For example, mouse movement and mouse button presses, optionally coordinated with single or multiple keyboard presses or holds; contact movements such as taps, drags, scrolls, etc. on touchpads; pen stylus inputs; movement of the device; oral instructions; detected eye movements; biometric inputs; and/or any combination thereof are optionally utilized as inputs corresponding to sub-events which define an event to be recognized.

FIG. **3** illustrates a portable multifunction device **200** having a touch screen **212** in accordance with some embodiments. The touch screen optionally displays one or more graphics within user interface (UI) **300**. In this embodiment, as well as others described below, a user is enabled to select one or more of the graphics by making a gesture on the

graphics, for example, with one or more fingers **302** (not drawn to scale in the figure) or one or more styluses **303** (not drawn to scale in the figure). In some embodiments, selection of one or more graphics occurs when the user breaks contact with the one or more graphics. In some embodiments, the gesture optionally includes one or more taps, one or more swipes (from left to right, right to left, upward and/or downward), and/or a rolling of a finger (from right to left, left to right, upward and/or downward) that has made contact with device **200**. In some implementations or circumstances, inadvertent contact with a graphic does not select the graphic. For example, a swipe gesture that sweeps over an application icon optionally does not select the corresponding application when the gesture corresponding to selection is a tap.

Device **200** also includes one or more physical buttons, such as “home” or menu button **304**. As described previously, menu button **304** is used to navigate to any application **236** in a set of applications that is executed on device **200**. Alternatively, in some embodiments, the menu button is implemented as a soft key in a GUI displayed on touch screen **212**.

In one embodiment, device **200** includes touch screen **212**, menu button **304**, push button **306** for powering the device on/off and locking the device, volume adjustment button(s) **308**, subscriber identity module (SIM) card slot **310**, headset jack **312**, and docking/charging external port **224**. Push button **306** is, optionally, used to turn the power on/off on the device by depressing the button and holding the button in the depressed state for a predefined time interval; to lock the device by depressing the button and releasing the button before the predefined time interval has elapsed; and/or to unlock the device or initiate an unlock process. In an alternative embodiment, device **200** also accepts verbal input for activation or deactivation of some functions through microphone **213**. Device **200** also, optionally, includes one or more contact intensity sensors **265** for detecting intensity of contacts on touch screen **212** and/or one or more tactile output generators **267** for generating tactile outputs for a user of device **200**.

FIG. 4 is a block diagram of an exemplary multifunction device with a display and a touch-sensitive surface in accordance with some embodiments. Device **400** need not be portable. In some embodiments, device **400** is a laptop computer, a desktop computer, a tablet computer, a multimedia player device, a navigation device, an educational device (such as a child’s learning toy), a gaming system, or a control device (e.g., a home or industrial controller). Device **400** typically includes one or more processing units (CPUs) **410**, one or more network or other communications interfaces **460**, memory **470**, and one or more communication buses **420** for interconnecting these components. Communication buses **420** optionally include circuitry (sometimes called a chipset) that interconnects and controls communications between system components. Device **400** includes input/output (I/O) interface **430** comprising display **440**, which is typically a touch screen display. I/O interface **430** also optionally includes a keyboard and/or mouse (or other pointing device) **450** and touchpad **455**, tactile output generator **457** for generating tactile outputs on device **400** (e.g., similar to tactile output generator(s) **267** described above with reference to FIG. 2A), sensors **459** (e.g., optical, acceleration, proximity, touch-sensitive, and/or contact intensity sensors similar to contact intensity sensor(s) **265** described above with reference to FIG. 2A). Memory **470** includes high-speed random access memory, such as DRAM, SRAM, DDR RAM, or other random access solid

state memory devices; and optionally includes non-volatile memory, such as one or more magnetic disk storage devices, optical disk storage devices, flash memory devices, or other non-volatile solid state storage devices. Memory **470** optionally includes one or more storage devices remotely located from CPU(s) **410**. In some embodiments, memory **470** stores programs, modules, and data structures analogous to the programs, modules, and data structures stored in memory **202** of portable multifunction device **200** (FIG. 2A), or a subset thereof. Furthermore, memory **470** optionally stores additional programs, modules, and data structures not present in memory **202** of portable multifunction device **200**. For example, memory **470** of device **400** optionally stores drawing module **480**, presentation module **482**, word processing module **484**, website creation module **486**, disk authoring module **488**, and/or spreadsheet module **490**, while memory **202** of portable multifunction device **200** (FIG. 2A) optionally does not store these modules.

Each of the above-identified elements in FIG. 4 is, in some examples, stored in one or more of the previously mentioned memory devices. Each of the above-identified modules corresponds to a set of instructions for performing a function described above. The above-identified modules or programs (e.g., sets of instructions) need not be implemented as separate software programs, procedures, or modules, and thus various subsets of these modules are combined or otherwise rearranged in various embodiments. In some embodiments, memory **470** stores a subset of the modules and data structures identified above. Furthermore, memory **470** stores additional modules and data structures not described above.

Attention is now directed towards embodiments of user interfaces that can be implemented on, for example, portable multifunction device **200**.

FIG. 5A illustrates an exemplary user interface for a menu of applications on portable multifunction device **200** in accordance with some embodiments. Similar user interfaces are implemented on device **400**. In some embodiments, user interface **500** includes the following elements, or a subset or superset thereof:

- Signal strength indicator(s) **502** for wireless communication(s), such as cellular and Wi-Fi signals;

- Time **504**;

- Bluetooth indicator **505**;

- Battery status indicator **506**;

- Tray **508** with icons for frequently used applications, such as:

- Icon **516** for telephone module **238**, labeled “Phone,” which optionally includes an indicator **514** of the number of missed calls or voicemail messages;

- Icon **518** for e-mail client module **240**, labeled “Mail,” which optionally includes an indicator **510** of the number of unread e-mails;

- Icon **520** for browser module **247**, labeled “Browser;” and

- Icon **522** for video and music player module **252**, also referred to as iPod (trademark of Apple Inc.) module **252**, labeled “iPod;” and

- Icons for other applications, such as:

- Icon **524** for IM module **241**, labeled “Messages;”

- Icon **526** for calendar module **248**, labeled “Calendar;”

- Icon **528** for image management module **244**, labeled “Photos;”

- Icon **530** for camera module **243**, labeled “Camera;”

- Icon **532** for online video module **255**, labeled “Online Video;”

- Icon **534** for stocks widget **249-2**, labeled “Stocks;”

Icon **536** for map module **254**, labeled “Maps;”
 Icon **538** for weather widget **249-1**, labeled “Weather;”
 Icon **540** for alarm clock widget **249-4**, labeled
 “Clock;”

Icon **542** for workout support module **242**, labeled
 “Workout Support;”

Icon **544** for notes module **253**, labeled “Notes;” and
 Icon **546** for a settings application or module, labeled
 “Settings;” which provides access to settings for
 device **200** and its various applications **236**.

It should be noted that the icon labels illustrated in FIG. **5A** are merely exemplary. For example, icon **522** for video and music player module **252** is optionally labeled “Music” or “Music Player.” Other labels are, optionally, used for various application icons. In some embodiments, a label for a respective application icon includes a name of an application corresponding to the respective application icon. In some embodiments, a label for a particular application icon is distinct from a name of an application corresponding to the particular application icon.

FIG. **5B** illustrates an exemplary user interface on a device (e.g., device **400**, FIG. **4**) with a touch-sensitive surface **551** (e.g., a tablet or touchpad **455**, FIG. **4**) that is separate from the display **550** (e.g., touch screen display **212**). Device **400** also, optionally, includes one or more contact intensity sensors (e.g., one or more of sensors **457**) for detecting intensity of contacts on touch-sensitive surface **551** and/or one or more tactile output generators **459** for generating tactile outputs for a user of device **400**.

Although some of the examples which follow will be given with reference to inputs on touch screen display **212** (where the touch-sensitive surface and the display are combined), in some embodiments, the device detects inputs on a touch-sensitive surface that is separate from the display, as shown in FIG. **5B**. In some embodiments, the touch-sensitive surface (e.g., **551** in FIG. **5B**) has a primary axis (e.g., **552** in FIG. **5B**) that corresponds to a primary axis (e.g., **553** in FIG. **5B**) on the display (e.g., **550**). In accordance with these embodiments, the device detects contacts (e.g., **560** and **562** in FIG. **5B**) with the touch-sensitive surface **551** at locations that correspond to respective locations on the display (e.g., in FIG. **5B**, **560** corresponds to **568** and **562** corresponds to **570**). In this way, user inputs (e.g., contacts **560** and **562**, and movements thereof) detected by the device on the touch-sensitive surface (e.g., **551** in FIG. **5B**) are used by the device to manipulate the user interface on the display (e.g., **550** in FIG. **5B**) of the multifunction device when the touch-sensitive surface is separate from the display. It should be understood that similar methods are, optionally, used for other user interfaces described herein.

Additionally, while the following examples are given primarily with reference to finger inputs (e.g., finger contacts, finger tap gestures, finger swipe gestures), it should be understood that, in some embodiments, one or more of the finger inputs are replaced with input from another input device (e.g., a mouse-based input or stylus input). For example, a swipe gesture is, optionally, replaced with a mouse click (e.g., instead of a contact) followed by movement of the cursor along the path of the swipe (e.g., instead of movement of the contact). As another example, a tap gesture is, optionally, replaced with a mouse click while the cursor is located over the location of the tap gesture (e.g., instead of detection of the contact followed by ceasing to detect the contact). Similarly, when multiple user inputs are simultaneously detected, it should be understood that mul-

tiple computer mice are, optionally, used simultaneously, or a mouse and finger contacts are, optionally, used simultaneously.

FIG. **6A** illustrates exemplary personal electronic device **600**. Device **600** includes body **602**. In some embodiments, device **600** includes some or all of the features described with respect to devices **200** and **400** (e.g., FIGS. **2A-4**). In some embodiments, device **600** has touch-sensitive display screen **604**, hereafter touch screen **604**. Alternatively, or in addition to touch screen **604**, device **600** has a display and a touch-sensitive surface. As with devices **200** and **400**, in some embodiments, touch screen **604** (or the touch-sensitive surface) has one or more intensity sensors for detecting intensity of contacts (e.g., touches) being applied. The one or more intensity sensors of touch screen **604** (or the touch-sensitive surface) provide output data that represents the intensity of touches. The user interface of device **600** responds to touches based on their intensity, meaning that touches of different intensities can invoke different user interface operations on device **600**.

Techniques for detecting and processing touch intensity are found, for example, in related applications: International Patent Application Serial No. PCT/US2013/040061, titled “Device, Method, and Graphical User Interface for Displaying User Interface Objects Corresponding to an Application,” filed May 8, 2013, and International Patent Application Serial No. PCT/US2013/069483, titled “Device, Method, and Graphical User Interface for Transitioning Between Touch Input to Display Output Relationships,” filed Nov. 11, 2013, each of which is hereby incorporated by reference in their entirety.

In some embodiments, device **600** has one or more input mechanisms **606** and **608**. Input mechanisms **606** and **608**, if included, are physical. Examples of physical input mechanisms include push buttons and rotatable mechanisms. In some embodiments, device **600** has one or more attachment mechanisms. Such attachment mechanisms, if included, can permit attachment of device **600** with, for example, hats, eyewear, earrings, necklaces, shirts, jackets, bracelets, watch straps, chains, trousers, belts, shoes, purses, backpacks, and so forth. These attachment mechanisms permit device **600** to be worn by a user.

FIG. **6B** depicts exemplary personal electronic device **600**. In some embodiments, device **600** includes some or all of the components described with respect to FIGS. **2A**, **2B**, and **4**. Device **600** has bus **612** that operatively couples I/O section **614** with one or more computer processors **616** and memory **618**. I/O section **614** is connected to display **604**, which can have touch-sensitive component **622** and, optionally, touch-intensity sensitive component **624**. In addition, I/O section **614** is connected with communication unit **630** for receiving application and operating system data, using Wi-Fi, Bluetooth, near field communication (NFC), cellular, and/or other wireless communication techniques. Device **600** includes input mechanisms **606** and/or **608**. Input mechanism **606** is a rotatable input device or a depressible and rotatable input device, for example. Input mechanism **608** is a button, in some examples.

Input mechanism **608** is a microphone, in some examples. Personal electronic device **600** includes, for example, various sensors, such as GPS sensor **632**, accelerometer **634**, directional sensor **640** (e.g., compass), gyroscope **636**, motion sensor **638**, and/or a combination thereof, all of which are operatively connected to I/O section **614**.

Memory **618** of personal electronic device **600** is a non-transitory computer-readable storage medium, for storing computer-executable instructions, which, when executed

by one or more computer processors 616, for example, cause the computer processors to perform the techniques and processes described below. The computer-executable instructions, for example, are also stored and/or transported within any non-transitory computer-readable storage medium for use by or in connection with an instruction execution system, apparatus, or device, such as a computer-based system, processor-containing system, or other system that can fetch the instructions from the instruction execution system, apparatus, or device and execute the instructions. Personal electronic device 600 is not limited to the components and configuration of FIG. 6B, but can include other or additional components in multiple configurations.

As used here, the term “affordance” refers to a user-interactive graphical user interface object that is, for example, displayed on the display screen of devices 200, 400, and/or 600 (FIGS. 2A, 4, and 6A-B). For example, an image (e.g., icon), a button, and text (e.g., hyperlink) each constitutes an affordance.

As used herein, the term “focus selector” refers to an input element that indicates a current part of a user interface with which a user is interacting. In some implementations that include a cursor or other location marker, the cursor acts as a “focus selector” so that when an input (e.g., a press input) is detected on a touch-sensitive surface (e.g., touchpad 455 in FIG. 4 or touch-sensitive surface 551 in FIG. 5B) while the cursor is over a particular user interface element (e.g., a button, window, slider or other user interface element), the particular user interface element is adjusted in accordance with the detected input. In some implementations that include a touch screen display (e.g., touch-sensitive display system 212 in FIG. 2A or touch screen 212 in FIG. 5A) that enables direct interaction with user interface elements on the touch screen display, a detected contact on the touch screen acts as a “focus selector” so that when an input (e.g., a press input by the contact) is detected on the touch screen display at a location of a particular user interface element (e.g., a button, window, slider, or other user interface element), the particular user interface element is adjusted in accordance with the detected input. In some implementations, focus is moved from one region of a user interface to another region of the user interface without corresponding movement of a cursor or movement of a contact on a touch screen display (e.g., by using a tab key or arrow keys to move focus from one button to another button); in these implementations, the focus selector moves in accordance with movement of focus between different regions of the user interface. Without regard to the specific form taken by the focus selector, the focus selector is generally the user interface element (or contact on a touch screen display) that is controlled by the user so as to communicate the user’s intended interaction with the user interface (e.g., by indicating, to the device, the element of the user interface with which the user is intending to interact). For example, the location of a focus selector (e.g., a cursor, a contact, or a selection box) over a respective button while a press input is detected on the touch-sensitive surface (e.g., a touchpad or touch screen) will indicate that the user is intending to activate the respective button (as opposed to other user interface elements shown on a display of the device).

As used in the specification and claims, the term “characteristic intensity” of a contact refers to a characteristic of the contact based on one or more intensities of the contact. In some embodiments, the characteristic intensity is based on multiple intensity samples. The characteristic intensity is, optionally, based on a predefined number of intensity samples, or a set of intensity samples collected during a

predetermined time period (e.g., 0.05, 0.1, 0.2, 0.5, 1, 2, 5, 10 seconds) relative to a predefined event (e.g., after detecting the contact, prior to detecting liftoff of the contact, before or after detecting a start of movement of the contact, prior to detecting an end of the contact, before or after detecting an increase in intensity of the contact, and/or before or after detecting a decrease in intensity of the contact). A characteristic intensity of a contact is, optionally based on one or more of: a maximum value of the intensities of the contact, a mean value of the intensities of the contact, an average value of the intensities of the contact, a top 10 percentile value of the intensities of the contact, a value at the half maximum of the intensities of the contact, a value at the 90 percent maximum of the intensities of the contact, or the like. In some embodiments, the duration of the contact is used in determining the characteristic intensity (e.g., when the characteristic intensity is an average of the intensity of the contact over time). In some embodiments, the characteristic intensity is compared to a set of one or more intensity thresholds to determine whether an operation has been performed by a user. For example, the set of one or more intensity thresholds includes a first intensity threshold and a second intensity threshold. In this example, a contact with a characteristic intensity that does not exceed the first threshold results in a first operation, a contact with a characteristic intensity that exceeds the first intensity threshold and does not exceed the second intensity threshold results in a second operation, and a contact with a characteristic intensity that exceeds the second threshold results in a third operation. In some embodiments, a comparison between the characteristic intensity and one or more thresholds is used to determine whether or not to perform one or more operations (e.g., whether to perform a respective operation or forgo performing the respective operation) rather than being used to determine whether to perform a first operation or a second operation.

In some embodiments, a portion of a gesture is identified for purposes of determining a characteristic intensity. For example, a touch-sensitive surface receives a continuous swipe contact transitioning from a start location and reaching an end location, at which point the intensity of the contact increases. In this example, the characteristic intensity of the contact at the end location is based on only a portion of the continuous swipe contact, and not the entire swipe contact (e.g., only the portion of the swipe contact at the end location). In some embodiments, a smoothing algorithm is applied to the intensities of the swipe contact prior to determining the characteristic intensity of the contact. For example, the smoothing algorithm optionally includes one or more of: an unweighted sliding-average smoothing algorithm, a triangular smoothing algorithm, a median filter smoothing algorithm, and/or an exponential smoothing algorithm. In some circumstances, these smoothing algorithms eliminate narrow spikes or dips in the intensities of the swipe contact for purposes of determining a characteristic intensity.

The intensity of a contact on the touch-sensitive surface is characterized relative to one or more intensity thresholds, such as a contact-detection intensity threshold, a light press intensity threshold, a deep press intensity threshold, and/or one or more other intensity thresholds. In some embodiments, the light press intensity threshold corresponds to an intensity at which the device will perform operations typically associated with clicking a button of a physical mouse or a trackpad. In some embodiments, the deep press intensity threshold corresponds to an intensity at which the device will perform operations that are different from operations

typically associated with clicking a button of a physical mouse or a trackpad. In some embodiments, when a contact is detected with a characteristic intensity below the light press intensity threshold (e.g., and above a nominal contact-detection intensity threshold below which the contact is no longer detected), the device will move a focus selector in accordance with movement of the contact on the touch-sensitive surface without performing an operation associated with the light press intensity threshold or the deep press intensity threshold. Generally, unless otherwise stated, these intensity thresholds are consistent between different sets of user interface figures.

An increase of characteristic intensity of the contact from an intensity below the light press intensity threshold to an intensity between the light press intensity threshold and the deep press intensity threshold is sometimes referred to as a “light press” input. An increase of characteristic intensity of the contact from an intensity below the deep press intensity threshold to an intensity above the deep press intensity threshold is sometimes referred to as a “deep press” input. An increase of characteristic intensity of the contact from an intensity below the contact-detection intensity threshold to an intensity between the contact-detection intensity threshold and the light press intensity threshold is sometimes referred to as detecting the contact on the touch-surface. A decrease of characteristic intensity of the contact from an intensity above the contact-detection intensity threshold to an intensity below the contact-detection intensity threshold is sometimes referred to as detecting liftoff of the contact from the touch-surface. In some embodiments, the contact-detection intensity threshold is zero. In some embodiments, the contact-detection intensity threshold is greater than zero.

In some embodiments described herein, one or more operations are performed in response to detecting a gesture that includes a respective press input or in response to detecting the respective press input performed with a respective contact (or a plurality of contacts), where the respective press input is detected based at least in part on detecting an increase in intensity of the contact (or plurality of contacts) above a press-input intensity threshold. In some embodiments, the respective operation is performed in response to detecting the increase in intensity of the respective contact above the press-input intensity threshold (e.g., a “down stroke” of the respective press input). In some embodiments, the press input includes an increase in intensity of the respective contact above the press-input intensity threshold and a subsequent decrease in intensity of the contact below the press-input intensity threshold, and the respective operation is performed in response to detecting the subsequent decrease in intensity of the respective contact below the press-input threshold (e.g., an “up stroke” of the respective press input).

In some embodiments, the device employs intensity hysteresis to avoid accidental inputs sometimes termed “jitter,” where the device defines or selects a hysteresis intensity threshold with a predefined relationship to the press-input intensity threshold (e.g., the hysteresis intensity threshold is X intensity units lower than the press-input intensity threshold or the hysteresis intensity threshold is 75%, 90%, or some reasonable proportion of the press-input intensity threshold). Thus, in some embodiments, the press input includes an increase in intensity of the respective contact above the press-input intensity threshold and a subsequent decrease in intensity of the contact below the hysteresis intensity threshold that corresponds to the press-input intensity threshold, and the respective operation is performed in response to detecting the subsequent decrease in intensity of

the respective contact below the hysteresis intensity threshold (e.g., an “up stroke” of the respective press input). Similarly, in some embodiments, the press input is detected only when the device detects an increase in intensity of the contact from an intensity at or below the hysteresis intensity threshold to an intensity at or above the press-input intensity threshold and, optionally, a subsequent decrease in intensity of the contact to an intensity at or below the hysteresis intensity, and the respective operation is performed in response to detecting the press input (e.g., the increase in intensity of the contact or the decrease in intensity of the contact, depending on the circumstances).

For ease of explanation, the descriptions of operations performed in response to a press input associated with a press-input intensity threshold or in response to a gesture including the press input are, optionally, triggered in response to detecting either: an increase in intensity of a contact above the press-input intensity threshold, an increase in intensity of a contact from an intensity below the hysteresis intensity threshold to an intensity above the press-input intensity threshold, a decrease in intensity of the contact below the press-input intensity threshold, and/or a decrease in intensity of the contact below the hysteresis intensity threshold corresponding to the press-input intensity threshold. Additionally, in examples where an operation is described as being performed in response to detecting a decrease in intensity of a contact below the press-input intensity threshold, the operation is, optionally, performed in response to detecting a decrease in intensity of the contact below a hysteresis intensity threshold corresponding to, and lower than, the press-input intensity threshold.

3. Digital Assistant System

FIG. 7A illustrates a block diagram of digital assistant system **700** in accordance with various examples. In some examples, digital assistant system **700** is implemented on a standalone computer system. In some examples, digital assistant system **700** is distributed across multiple computers. In some examples, some of the modules and functions of the digital assistant are divided into a server portion and a client portion, where the client portion resides on one or more user devices (e.g., devices **104**, **122**, **200**, **400**, or **600**) and communicates with the server portion (e.g., server system **108**) through one or more networks, e.g., as shown in FIG. 1. In some examples, digital assistant system **700** is an implementation of server system **108** (and/or DA server **106**) shown in FIG. 1. It should be noted that digital assistant system **700** is only one example of a digital assistant system, and that digital assistant system **700** can have more or fewer components than shown, can combine two or more components, or can have a different configuration or arrangement of the components. The various components shown in FIG. 7A are implemented in hardware, software instructions for execution by one or more processors, firmware, including one or more signal processing and/or application specific integrated circuits, or a combination thereof.

Digital assistant system **700** includes memory **702**, one or more processors **704**, input/output (I/O) interface **706**, and network communications interface **708**. These components can communicate with one another over one or more communication buses or signal lines **710**.

In some examples, memory **702** includes a non-transitory computer-readable medium, such as high-speed random access memory and/or a non-volatile computer-readable storage medium (e.g., one or more magnetic disk storage devices, flash memory devices, or other non-volatile solid-state memory devices).

In some examples, I/O interface **706** couples input/output devices **716** of digital assistant system **700**, such as displays, keyboards, touch screens, and microphones, to user interface module **722**. I/O interface **706**, in conjunction with user interface module **722**, receives user inputs (e.g., voice input, keyboard inputs, touch inputs, etc.) and processes them accordingly. In some examples, e.g., when the digital assistant is implemented on a standalone user device, digital assistant system **700** includes any of the components and I/O communication interfaces described with respect to devices **200**, **400**, or **600** in FIGS. **2A**, **4**, **6A-B**, respectively. In some examples, digital assistant system **700** represents the server portion of a digital assistant implementation, and can interact with the user through a client-side portion residing on a user device (e.g., devices **104**, **200**, **400**, or **600**).

In some examples, the network communications interface **708** includes wired communication port(s) **712** and/or wireless transmission and reception circuitry **714**. The wired communication port(s) receives and send communication signals via one or more wired interfaces, e.g., Ethernet, Universal Serial Bus (USB), FIREWIRE, etc. The wireless circuitry **714** receives and sends RF signals and/or optical signals from/to communications networks and other communications devices. The wireless communications use any of a plurality of communications standards, protocols, and technologies, such as GSM, EDGE, CDMA, TDMA, Bluetooth, Wi-Fi, VoIP, Wi-MAX, or any other suitable communication protocol. Network communications interface **708** enables communication between digital assistant system **700** with networks, such as the Internet, an intranet, and/or a wireless network, such as a cellular telephone network, a wireless local area network (LAN), and/or a metropolitan area network (MAN), and other devices.

In some examples, memory **702**, or the computer-readable storage media of memory **702**, stores programs, modules, instructions, and data structures including all or a subset of: operating system **718**, communications module **720**, user interface module **722**, one or more applications **724**, and digital assistant module **726**. In particular, memory **702**, or the computer-readable storage media of memory **702**, stores instructions for performing the processes described below. One or more processors **704** execute these programs, modules, and instructions, and reads/writes from/to the data structures.

Operating system **718** (e.g., Darwin, RTXC, LINUX, UNIX, iOS, OS X, WINDOWS, or an embedded operating system such as VxWorks) includes various software components and/or drivers for controlling and managing general system tasks (e.g., memory management, storage device control, power management, etc.) and facilitates communications between various hardware, firmware, and software components.

Communications module **720** facilitates communications between digital assistant system **700** with other devices over network communications interface **708**. For example, communications module **720** communicates with RF circuitry **208** of electronic devices such as devices **200**, **400**, and **600** shown in FIGS. **2A**, **4**, **6A-B**, respectively. Communications module **720** also includes various components for handling data received by wireless circuitry **714** and/or wired communications port **712**.

User interface module **722** receives commands and/or inputs from a user via I/O interface **706** (e.g., from a keyboard, touch screen, pointing device, controller, and/or microphone), and generate user interface objects on a display. User interface module **722** also prepares and delivers outputs (e.g., speech, sound, animation, text, icons, vibra-

tions, haptic feedback, light, etc.) to the user via the I/O interface **706** (e.g., through displays, audio channels, speakers, touch-pads, etc.).

Applications **724** include programs and/or modules that are configured to be executed by one or more processors **704**. For example, if the digital assistant system is implemented on a standalone user device, applications **724** include user applications, such as games, a calendar application, a navigation application, or an email application. If digital assistant system **700** is implemented on a server, applications **724** include resource management applications, diagnostic applications, or scheduling applications, for example.

Memory **702** also stores digital assistant module **726** (or the server portion of a digital assistant). In some examples, digital assistant module **726** includes the following sub-modules, or a subset or superset thereof: input/output processing module **728**, speech-to-text (STT) processing module **730**, natural language processing module **732**, dialogue flow processing module **734**, task flow processing module **736**, service processing module **738**, and speech synthesis processing module **740**. Each of these modules has access to one or more of the following systems or data and models of the digital assistant module **726**, or a subset or superset thereof: ontology **760**, vocabulary index **744**, user data **748**, task flow models **754**, service models **756**, and ASR systems **758**.

In some examples, using the processing modules, data, and models implemented in digital assistant module **726**, the digital assistant can perform at least some of the following: converting speech input into text; identifying a user's intent expressed in a natural language input received from the user; actively eliciting and obtaining information needed to fully infer the user's intent (e.g., by disambiguating words, games, intentions, etc.); determining the task flow for fulfilling the inferred intent; and executing the task flow to fulfill the inferred intent.

In some examples, as shown in FIG. **7B**, I/O processing module **728** interacts with the user through I/O devices **716** in FIG. **7A** or with a user device (e.g., devices **104**, **200**, **400**, or **600**) through network communications interface **708** in FIG. **7A** to obtain user input (e.g., a speech input) and to provide responses (e.g., as speech outputs) to the user input. I/O processing module **728** optionally obtains contextual information associated with the user input from the user device, along with or shortly after the receipt of the user input. The contextual information includes user-specific data, vocabulary, and/or preferences relevant to the user input. In some examples, the contextual information also includes software and hardware states of the user device at the time the user request is received, and/or information related to the surrounding environment of the user at the time that the user request was received. In some examples, I/O processing module **728** also sends follow-up questions to, and receive answers from, the user regarding the user request. When a user request is received by I/O processing module **728** and the user request includes speech input, I/O processing module **728** forwards the speech input to STT processing module **730** (or speech recognizer) for speech-to-text conversions.

STT processing module **730** includes one or more ASR systems **758**. The one or more ASR systems **758** can process the speech input that is received through I/O processing module **728** to produce a recognition result. Each ASR system **758** includes a front-end speech pre-processor. The front-end speech pre-processor extracts representative features from the speech input. For example, the front-end

speech pre-processor performs a Fourier transform on the speech input to extract spectral features that characterize the speech input as a sequence of representative multi-dimensional vectors. Further, each ASR system 758 includes one or more speech recognition models (e.g., acoustic models and/or language models) and implements one or more speech recognition engines. Examples of speech recognition models include Hidden Markov Models, Gaussian-Mixture Models, Deep Neural Network Models, n-gram language models, and other statistical models. Examples of speech recognition engines include the dynamic time warping based engines and weighted finite-state transducers (WFST) based engines. The one or more speech recognition models and the one or more speech recognition engines are used to process the extracted representative features of the front-end speech pre-processor to produce intermediate recognitions results (e.g., phonemes, phonemic strings, and sub-words), and ultimately, text recognition results (e.g., words, word strings, or sequence of tokens). In some examples, the speech input is processed at least partially by a third-party service or on the user's device (e.g., device 104, 200, 400, or 600) to produce the recognition result. Once STT processing module 730 produces recognition results containing a text string (e.g., words, or sequence of words, or sequence of tokens), the recognition result is passed to natural language processing module 732 for intent deduction. In some examples, STT processing module 730 produces multiple candidate text representations of the speech input. Each candidate text representation is a sequence of words or tokens corresponding to the speech input. In some examples, each candidate text representation is associated with a speech recognition confidence score. Based on the speech recognition confidence scores, STT processing module 730 ranks the candidate text representations and provides the n-best (e.g., n highest ranked) candidate text representation(s) to natural language processing module 732 for intent deduction, where n is a predetermined integer greater than zero. For example, in one example, only the highest ranked (n=1) candidate text representation is passed to natural language processing module 732 for intent deduction. In another example, the five highest ranked (n=5) candidate text representations are passed to natural language processing module 732 for intent deduction.

More details on the speech-to-text processing are described in U.S. Utility application Ser. No. 13/236,942 for "Consolidating Speech Recognition Results," filed on Sep. 20, 2011, the entire disclosure of which is incorporated herein by reference.

In some examples, STT processing module 730 includes and/or accesses a vocabulary of recognizable words via phonetic alphabet conversion module 731. Each vocabulary word is associated with one or more candidate pronunciations of the word represented in a speech recognition phonetic alphabet. In particular, the vocabulary of recognizable words includes a word that is associated with a plurality of candidate pronunciations. For example, the vocabulary includes the word "tomato" that is associated with the candidate pronunciations of /tə'meɪroʊ/ and /tə'matoʊ/. Further, vocabulary words are associated with custom candidate pronunciations that are based on previous speech inputs from the user. Such custom candidate pronunciations are stored in STT processing module 730 and are associated with a particular user via the user's profile on the device. In some examples, the candidate pronunciations for words are determined based on the spelling of the word and one or more linguistic and/or phonetic rules. In some examples, the

candidate pronunciations are manually generated, e.g., based on known canonical pronunciations.

In some examples, the candidate pronunciations are ranked based on the commonness of the candidate pronunciation. For example, the candidate pronunciation /tə'meɪroʊ/ is ranked higher than /tə'matoʊ/, because the former is a more commonly used pronunciation (e.g., among all users, for users in a particular geographical region, or for any other appropriate subset of users). In some examples, candidate pronunciations are ranked based on whether the candidate pronunciation is a custom candidate pronunciation associated with the user. For example, custom candidate pronunciations are ranked higher than canonical candidate pronunciations. This can be useful for recognizing proper nouns having a unique pronunciation that deviates from canonical pronunciation. In some examples, candidate pronunciations are associated with one or more speech characteristics, such as geographic origin, nationality, or ethnicity. For example, the candidate pronunciation /tə'meɪroʊ/ is associated with the United States, whereas the candidate pronunciation /tə'matoʊ/ is associated with Great Britain. Further, the rank of the candidate pronunciation is based on one or more characteristics (e.g., geographic origin, nationality, ethnicity, etc.) of the user stored in the user's profile on the device. For example, it can be determined from the user's profile that the user is associated with the United States. Based on the user being associated with the United States, the candidate pronunciation /tə'meɪroʊ/ (associated with the United States) is ranked higher than the candidate pronunciation /tə'matoʊ/ (associated with Great Britain). In some examples, one of the ranked candidate pronunciations is selected as a predicted pronunciation (e.g., the most likely pronunciation).

When a speech input is received, STT processing module 730 is used to determine the phonemes corresponding to the speech input (e.g., using an acoustic model), and then attempt to determine words that match the phonemes (e.g., using a language model). For example, if STT processing module 730 first identifies the sequence of phonemes /tə'meɪroʊ/ corresponding to a portion of the speech input, it can then determine, based on vocabulary index 744, that this sequence corresponds to the word "tomato."

In some examples, STT processing module 730 uses approximate matching techniques to determine words in an utterance. Thus, for example, the STT processing module 730 determines that the sequence of phonemes /tə'meɪroʊ/ corresponds to the word "tomato," even if that particular sequence of phonemes is not one of the candidate sequence of phonemes for that word.

Natural language processing module 732 ("natural language processor") of the digital assistant takes the n-best candidate text representation(s) ("word sequence(s)" or "token sequence(s)") generated by STT processing module 730, and attempts to associate each of the candidate text representations with one or more "actionable intents" recognized by the digital assistant. An "actionable intent" (or "user intent") represents a task that can be performed by the digital assistant, and can have an associated task flow implemented in task flow models 754. The associated task flow is a series of programmed actions and steps that the digital assistant takes in order to perform the task. The scope of a digital assistant's capabilities is dependent on the number and variety of task flows that have been implemented and stored in task flow models 754, or in other words, on the number and variety of "actionable intents" that the digital assistant recognizes. The effectiveness of the digital assistant, however, also depends on the assistant's

ability to infer the correct “actionable intent(s)” from the user request expressed in natural language.

In some examples, in addition to the sequence of words or tokens obtained from STT processing module 730, natural language processing module 732 also receives contextual information associated with the user request, e.g., from I/O processing module 728. The natural language processing module 732 optionally uses the contextual information to clarify, supplement, and/or further define the information contained in the candidate text representations received from STT processing module 730. The contextual information includes, for example, user preferences, hardware, and/or software states of the user device, sensor information collected before, during, or shortly after the user request, prior interactions (e.g., dialogue) between the digital assistant and the user, and the like. As described herein, contextual information is, in some examples, dynamic, and changes with time, location, content of the dialogue, and other factors.

In some examples, the natural language processing is based on, e.g., ontology 760. Ontology 760 is a hierarchical structure containing many nodes, each node representing either an “actionable intent” or a “property” relevant to one or more of the “actionable intents” or other “properties.” As noted above, an “actionable intent” represents a task that the digital assistant is capable of performing, i.e., it is “actionable” or can be acted on. A “property” represents a parameter associated with an actionable intent or a sub-aspect of another property. A linkage between an actionable intent node and a property node in ontology 760 defines how a parameter represented by the property node pertains to the task represented by the actionable intent node.

In some examples, ontology 760 is made up of actionable intent nodes and property nodes. Within ontology 760, each actionable intent node is linked to one or more property nodes either directly or through one or more intermediate property nodes. Similarly, each property node is linked to one or more actionable intent nodes either directly or through one or more intermediate property nodes. For example, as shown in FIG. 7C, ontology 760 includes a “restaurant reservation” node (i.e., an actionable intent node). Property nodes “restaurant,” “date/time” (for the reservation), and “party size” are each directly linked to the actionable intent node (i.e., the “restaurant reservation” node).

In addition, property nodes “cuisine,” “price range,” “phone number,” and “location” are sub-nodes of the property node “restaurant,” and are each linked to the “restaurant reservation” node (i.e., the actionable intent node) through the intermediate property node “restaurant.” For another example, as shown in FIG. 7C, ontology 760 also includes a “set reminder” node (i.e., another actionable intent node). Property nodes “date/time” (for setting the reminder) and “subject” (for the reminder) are each linked to the “set reminder” node. Since the property “date/time” is relevant to both the task of making a restaurant reservation and the task of setting a reminder, the property node “date/time” is linked to both the “restaurant reservation” node and the “set reminder” node in ontology 760.

An actionable intent node, along with its linked property nodes, is described as a “domain.” In the present discussion, each domain is associated with a respective actionable intent, and refers to the group of nodes (and the relationships there between) associated with the particular actionable intent. For example, ontology 760 shown in FIG. 7C includes an example of restaurant reservation domain 762 and an example of reminder domain 764 within ontology

760. The restaurant reservation domain includes the actionable intent node “restaurant reservation,” property nodes “restaurant,” “date/time,” and “party size,” and sub-property nodes “cuisine,” “price range,” “phone number,” and “location.” Reminder domain 764 includes the actionable intent node “set reminder,” and property nodes “subject” and “date/time.” In some examples, ontology 760 is made up of many domains. Each domain shares one or more property nodes with one or more other domains. For example, the “date/time” property node is associated with many different domains (e.g., a scheduling domain, a travel reservation domain, a movie ticket domain, etc.), in addition to restaurant reservation domain 762 and reminder domain 764.

While FIG. 7C illustrates two example domains within ontology 760, other domains include, for example, “find a movie,” “initiate a phone call,” “find directions,” “schedule a meeting,” “send a message,” and “provide an answer to a question,” “read a list,” “providing navigation instructions,” “provide instructions for a task” and so on. A “send a message” domain is associated with a “send a message” actionable intent node, and further includes property nodes such as “recipient(s),” “message type,” and “message body.” The property node “recipient” is further defined, for example, by the sub-property nodes such as “recipient name” and “message address.”

In some examples, ontology 760 includes all the domains (and hence actionable intents) that the digital assistant is capable of understanding and acting upon. In some examples, ontology 760 is modified, such as by adding or removing entire domains or nodes, or by modifying relationships between the nodes within the ontology 760.

In some examples, nodes associated with multiple related actionable intents are clustered under a “super domain” in ontology 760. For example, a “travel” super-domain includes a cluster of property nodes and actionable intent nodes related to travel. The actionable intent nodes related to travel includes “airline reservation,” “hotel reservation,” “car rental,” “get directions,” “find points of interest,” and so on. The actionable intent nodes under the same super domain (e.g., the “travel” super domain) have many property nodes in common. For example, the actionable intent nodes for “airline reservation,” “hotel reservation,” “car rental,” “get directions,” and “find points of interest” share one or more of the property nodes “start location,” “destination,” “departure date/time,” “arrival date/time,” and “party size.”

In some examples, each node in ontology 760 is associated with a set of words and/or phrases that are relevant to the property or actionable intent represented by the node. The respective set of words and/or phrases associated with each node are the so-called “vocabulary” associated with the node. The respective set of words and/or phrases associated with each node are stored in vocabulary index 744 in association with the property or actionable intent represented by the node. For example, returning to FIG. 7B, the vocabulary associated with the node for the property of “restaurant” includes words such as “food,” “drinks,” “cuisine,” “hungry,” “eat,” “pizza,” “fast food,” “meal,” and so on. For another example, the vocabulary associated with the node for the actionable intent of “initiate a phone call” includes words and phrases such as “call,” “phone,” “dial,” “ring,” “call this number,” “make a call to,” and so on. The vocabulary index 744 optionally includes words and phrases in different languages.

Natural language processing module 732 receives the candidate text representations (e.g., text string(s) or token sequence(s)) from STT processing module 730, and for each

candidate representation, determines what nodes are implicated by the words in the candidate text representation. In some examples, if a word or phrase in the candidate text representation is found to be associated with one or more nodes in ontology **760** (via vocabulary index **744**), the word or phrase “triggers” or “activates” those nodes. Based on the quantity and/or relative importance of the activated nodes, natural language processing module **732** selects one of the actionable intents as the task that the user intended the digital assistant to perform. In some examples, the domain that has the most “triggered” nodes is selected. In some examples, the domain having the highest confidence value (e.g., based on the relative importance of its various triggered nodes) is selected. In some examples, the domain is selected based on a combination of the number and the importance of the triggered nodes. In some examples, additional factors are considered in selecting the node as well, such as whether the digital assistant has previously correctly interpreted a similar request from a user.

User data **748** includes user-specific information, such as user-specific vocabulary, user preferences, user address, user’s default and secondary languages, user’s contact list, and other short-term or long-term information for each user. In some examples, natural language processing module **732** uses the user-specific information to supplement the information contained in the user input to further define the user intent. For example, for a user request “invite my friends to my birthday party,” natural language processing module **732** is able to access user data **748** to determine who the “friends” are and when and where the “birthday party” would be held, rather than requiring the user to provide such information explicitly in his/her request.

It should be recognized that in some examples, natural language processing module **732** is implemented using one or more machine learning mechanisms (e.g., neural networks). In particular, the one or more machine learning mechanisms are configured to receive a candidate text representation and contextual information associated with the candidate text representation. Based on the candidate text representation and the associated contextual information, the one or more machine learning mechanisms are configured to determine intent confidence scores over a set of candidate actionable intents. Natural language processing module **732** can select one or more candidate actionable intents from the set of candidate actionable intents based on the determined intent confidence scores. In some examples, an ontology (e.g., ontology **760**) is also used to select the one or more candidate actionable intents from the set of candidate actionable intents.

Other details of searching an ontology based on a token string are described in U.S. Utility application Ser. No. 12/341,743 for “Method and Apparatus for Searching Using An Active Ontology,” filed Dec. 22, 2008, the entire disclosure of which is incorporated herein by reference.

In some examples, once natural language processing module **732** identifies an actionable intent (or domain) based on the user request, natural language processing module **732** generates a structured query to represent the identified actionable intent. In some examples, the structured query includes parameters for one or more nodes within the domain for the actionable intent, and at least some of the parameters are populated with the specific information and requirements specified in the user request. For example, the user says “Make me a dinner reservation at a sushi place at 7.” In this case, natural language processing module **732** is able to correctly identify the actionable intent to be “restaurant reservation” based on the user input. According to the

ontology, a structured query for a “restaurant reservation” domain includes parameters such as {Cuisine}, {Time}, {Date}, {Party Size}, and the like. In some examples, based on the speech input and the text derived from the speech input using STT processing module **730**, natural language processing module **732** generates a partial structured query for the restaurant reservation domain, where the partial structured query includes the parameters {Cuisine=“Sushi”} and {Time=“7 pm”}. However, in this example, the user’s utterance contains insufficient information to complete the structured query associated with the domain. Therefore, other necessary parameters such as {Party Size} and {Date} are not specified in the structured query based on the information currently available. In some examples, natural language processing module **732** populates some parameters of the structured query with received contextual information. For example, in some examples, if the user requested a sushi restaurant “near me,” natural language processing module **732** populates a {location} parameter in the structured query with GPS coordinates from the user device.

In some examples, natural language processing module **732** identifies multiple candidate actionable intents for each candidate text representation received from STT processing module **730**. Further, in some examples, a respective structured query (partial or complete) is generated for each identified candidate actionable intent. Natural language processing module **732** determines an intent confidence score for each candidate actionable intent and ranks the candidate actionable intents based on the intent confidence scores. In some examples, natural language processing module **732** passes the generated structured query (or queries), including any completed parameters, to task flow processing module **736** (“task flow processor”). In some examples, the structured query (or queries) for the m-best (e.g., m highest ranked) candidate actionable intents are provided to task flow processing module **736**, where m is a predetermined integer greater than zero. In some examples, the structured query (or queries) for the m-best candidate actionable intents are provided to task flow processing module **736** with the corresponding candidate text representation(s).

Other details of inferring a user intent based on multiple candidate actionable intents determined from multiple candidate text representations of a speech input are described in U.S. Utility application Ser. No. 14/298,725 for “System and Method for Inferring User Intent From Speech Inputs,” filed Jun. 6, 2014, the entire disclosure of which is incorporated herein by reference.

Task flow processing module **736** is configured to receive the structured query (or queries) from natural language processing module **732**, complete the structured query, if necessary, and perform the actions required to “complete” the user’s ultimate request. In some examples, the various procedures necessary to complete these tasks are provided in task flow models **754**. In some examples, task flow models **754** include procedures for obtaining additional information from the user and task flows for performing actions associated with the actionable intent.

As described above, in order to complete a structured query, task flow processing module **736** needs to initiate additional dialogue with the user in order to obtain additional information, and/or disambiguate potentially ambiguous utterances. When such interactions are necessary, task flow processing module **736** invokes dialogue flow processing module **734** to engage in a dialogue with the user. In some examples, dialogue flow processing module **734** determines how (and/or when) to ask the user for the additional information and receives and processes the user responses.

The questions are provided to and answers are received from the users through I/O processing module 728. In some examples, dialogue flow processing module 734 presents dialogue output to the user via audio and/or visual output, and receives input from the user via spoken or physical (e.g., clicking) responses. Continuing with the example above, when task flow processing module 736 invokes dialogue flow processing module 734 to determine the “party size” and “date” information for the structured query associated with the domain “restaurant reservation,” dialogue flow processing module 734 generates questions such as “For how many people?” and “On which day?” to pass to the user. Once answers are received from the user, dialogue flow processing module 734 then populates the structured query with the missing information, or pass the information to task flow processing module 736 to complete the missing information from the structured query.

Once task flow processing module 736 has completed the structured query for an actionable intent, task flow processing module 736 proceeds to perform the ultimate task associated with the actionable intent. Accordingly, task flow processing module 736 executes the steps and instructions in the task flow model according to the specific parameters contained in the structured query. For example, the task flow model for the actionable intent of “restaurant reservation” includes steps and instructions for contacting a restaurant and actually requesting a reservation for a particular party size at a particular time. For example, using a structured query such as: {restaurant reservation, restaurant=ABC Café, date=3/12/2012, time=7 pm, party size=5}, task flow processing module 736 performs the steps of: (1) logging onto a server of the ABC Café or a restaurant reservation system such as OPENTABLE®, (2) entering the date, time, and party size information in a form on the website, (3) submitting the form, and (4) making a calendar entry for the reservation in the user’s calendar.

In some examples, task flow processing module 736 employs the assistance of service processing module 738 (“service processing module”) to complete a task requested in the user input or to provide an informational answer requested in the user input. For example, service processing module 738 acts on behalf of task flow processing module 736 to make a phone call, set a calendar entry, invoke a map search, invoke or interact with other user applications installed on the user device, and invoke or interact with third-party services (e.g., a restaurant reservation portal, a social networking website, a banking portal, etc.). In some examples, the protocols and application programming interfaces (API) required by each service are specified by a respective service model among service models 756. Service processing module 738 accesses the appropriate service model for a service and generates requests for the service in accordance with the protocols and APIs required by the service according to the service model.

For example, if a restaurant has enabled an online reservation service, the restaurant submits a service model specifying the necessary parameters for making a reservation and the APIs for communicating the values of the necessary parameter to the online reservation service. When requested by task flow processing module 736, service processing module 738 establishes a network connection with the online reservation service using the web address stored in the service model, and sends the necessary parameters of the reservation (e.g., time, date, party size) to the online reservation interface in a format according to the API of the online reservation service.

In some examples, natural language processing module 732, dialogue flow processing module 734, and task flow processing module 736 are used collectively and iteratively to infer and define the user’s intent, obtain information to further clarify and refine the user intent, and finally generate a response (i.e., an output to the user, or the completion of a task) to fulfill the user’s intent. The generated response is a dialogue response to the speech input that at least partially fulfills the user’s intent. Further, in some examples, the generated response is output as a speech output. In these examples, the generated response is sent to speech synthesis processing module 740 (e.g., speech synthesizer) where it can be processed to synthesize the dialogue response in speech form. In yet other examples, the generated response is data content relevant to satisfying a user request in the speech input.

In examples where task flow processing module 736 receives multiple structured queries from natural language processing module 732, task flow processing module 736 initially processes the first structured query of the received structured queries to attempt to complete the first structured query and/or execute one or more tasks or actions represented by the first structured query. In some examples, the first structured query corresponds to the highest ranked actionable intent. In other examples, the first structured query is selected from the received structured queries based on a combination of the corresponding speech recognition confidence scores and the corresponding intent confidence scores. In some examples, if task flow processing module 736 encounters an error during processing of the first structured query (e.g., due to an inability to determine a necessary parameter), the task flow processing module 736 can proceed to select and process a second structured query of the received structured queries that corresponds to a lower ranked actionable intent. The second structured query is selected, for example, based on the speech recognition confidence score of the corresponding candidate text representation, the intent confidence score of the corresponding candidate actionable intent, a missing necessary parameter in the first structured query, or any combination thereof.

Speech synthesis processing module 740 is configured to synthesize speech outputs for presentation to the user. Speech synthesis processing module 740 synthesizes speech outputs based on text provided by the digital assistant. For example, the generated dialogue response is in the form of a text string. Speech synthesis processing module 740 converts the text string to an audible speech output. Speech synthesis processing module 740 uses any appropriate speech synthesis technique in order to generate speech outputs from text, including, but not limited to, concatenative synthesis, unit selection synthesis, diphone synthesis, domain-specific synthesis, formant synthesis, articulatory synthesis, hidden Markov model (HMM) based synthesis, and sinewave synthesis. In some examples, speech synthesis processing module 740 is configured to synthesize individual words based on phonemic strings corresponding to the words. For example, a phonemic string is associated with a word in the generated dialogue response. The phonemic string is stored in metadata associated with the word. Speech synthesis processing module 740 is configured to directly process the phonemic string in the metadata to synthesize the word in speech form.

In some examples, instead of (or in addition to) using speech synthesis processing module 740, speech synthesis is performed on a remote device (e.g., the server system 108), and the synthesized speech is sent to the user device for output to the user. For example, this can occur in some

implementations where outputs for a digital assistant are generated at a server system. And because server systems generally have more processing power or resources than a user device, it is possible to obtain higher quality speech outputs than would be practical with client-side synthesis.

Additional details on digital assistants can be found in the U.S. Utility application Ser. No. 12/987,982, entitled "Intelligent Automated Assistant," filed Jan. 10, 2011, and U.S. Utility application Ser. No. 13/251,088, entitled "Generating and Processing Task Items That Represent Tasks to Perform," filed Sep. 30, 2011, the entire disclosures of which are incorporated herein by reference.

4. Systems and Processes for Creating and Updating a Natural Language Model

FIG. 8 illustrates a block diagram of a portion of system 800 for creating a natural language model for an application, according to various examples. In some examples, system 800 is implemented on one or more electronic devices (e.g., devices 104, 122, 200, 400, or 600) and the modules and functions of system 800 may be distributed in any manner between the devices. In some examples, some of the modules and functions of system 800 are divided into a server portion and a client portion, where the client portion resides on one or more user devices (e.g., devices 104, 122, 200, 400, and 600) and communicates with the server portion (e.g., server system 108) through one or more networks, e.g., as shown in FIG. 1. System 800 is implemented using hardware, software, or a combination of hardware and software to carry out the principles discussed herein.

In some examples, the sub-modules or a subset or superset of the sub-modules described with regard to system 800 may be implemented in a digital assistant such as digital assistant system 700, discussed above. In particular, as is discussed in more detail, system 800 creates a natural language model for one or more applications which may be integrated with a digital assistant system, enabling the digital assistant system to determine a user intent and/or a task from a user input using natural language processing. Further, it will be understood that in some examples, system 800 may include one or more of the components discussed above including components for automatic speech recognition or speech to text capabilities.

It should be noted that system 800 is exemplary, and thus system 800 can have more or fewer components than shown, can combine two or more components, or can have a different configuration or arrangement of the components. Further, although the below discussion describes functions being performed at a single component of system 800, it is to be understood that such functions can be performed at other components of system 800 and that such functions can be performed at more than one component of system 800.

As illustrated in FIG. 8, system 800 includes natural language model generator 804 which performs a plurality of functions including action structure association 806, utterance augmentation 810, and utterance mapping 814. System 800 receives one or more example utterances 802 and provides them to natural language model generator 804 to create natural language model 816 using action structure association 806, utterance augmentation 810, and utterance mapping 814 as described further below.

In particular, system 800 may generate natural language model 816 for one or more applications that are to be integrated with a digital assistant system, electronic device, or other system. Accordingly, system 800 generates models for recognizing user intents and tasks from user input based on example utterances 802. In some examples, utterances 802 include training data for natural language model gen-

erator 804 to create natural language model 816. In some examples, utterances 802 are part of a training data for natural language model generator 804 to create natural language model 816. In some examples, utterances 802 and natural language model 816 are associated with a particular application.

In some examples, utterances 802 include example utterances that a developer of an application believes the application will receive and should respond to. For example, as shown in FIG. 8, utterances 802 include the example utterances "Get me a ride to SFO," and "Where is my car?." Accordingly, utterances 802 may be provided by a developer of a ride share application who wishes the application to recognize and respond to requests related to calling a ride and providing the status of the ride.

In some examples, utterances 802 are provided to system 800 in response to a prompt for example utterances. For example, a developer may interact with system 800 and be provided with a prompt to provide training data. Accordingly, the developer may provide training data that can include utterances 802 designated as example utterances. In some examples, utterances 802 are provided as a textual input (e.g., a file including text of example utterances). In some examples, utterances 802 are provided as a spoken input. For example, during an interaction with system 800, system 800 may provide a spoken output or other prompt for example utterances. A developer of the application or other user may respond by speaking utterances 802 to system 800. System 800 may then convert the spoken utterances to text using speech to text components (not pictured) as discussed above.

After receiving utterances 802, system 800 provides utterances 802 to natural language model generator 804 which performs action structure association 806. Action structure association 806 includes the process of associating action structure 808 with utterance 802.

In some examples, action structure 808 is provided with utterance 802. Further, the association of action structure 808 with utterance 802 may be received by system 800 at the same time from the user or developer. For example, when utterances 802 include "Get me a ride to SFO," system 800 may receive action structure 808 as shown in FIG. 8. Thus, system 800 receives "Get me a ride to SFO" along with the action structure "Rideshare: [end location];" and a mapping of the term "SFO" to the parameter [end location].

In some examples, action structure 808 includes a template with one or more parameters. Further, the parameter of the template is a parameter to be used when determining and performing a task determined based on the provided utterance associated with the action structure. For example, as discussed above, a received action structure may be "Rideshare: [end location]" in which the [end location] is a parameter of the template "Rideshare: [end location]." In particular, this action structure specifies the task of getting a rideshare and an associated parameter required of an ending location. Accordingly, when the utterance "Get me a ride to SFO," is received, natural language model generator 804 recognizes that "SFO" is a value for the parameter "end location" of the associated action structure which can be used to execute the task of booking a ride share.

In some examples, action structure 808 (or other action structures of system 800) is represented by a data structure. For example, action structure 808 may be stored or represented as a data structure including a task and one or more parameters such as the task "rideshare" and the parameter "end location." In some examples, action structure 808 can be represented by any known structure for storing and

retrieving data in an accessible format. Accordingly, action structure **808** may be stored and referenced by using a data structure so that action structure **808** is easily accessible to system **800**, other electronic devices, or a digital assistant as discussed further below.

In some examples, the parameter of action structure **808** is associated with the application corresponding to natural language model **816**. For example, if natural language model **816** is being created for a rideshare application, as previously discussed, then parameters of action structure **808** may be “starting location,” “type of vehicle,” “end location,” or other parameters typically associated with requesting or providing a ride share.

In some examples, the parameter of action structure **808** is associated with the application corresponding to natural language model **816** and one or more other applications. For example, as discussed above, natural language model **816** may be created for a rideshare application and parameters of action structure **808** may be associated with getting a rideshare. However, the parameters of action structure **808** may also be associated with an application for booking a flight or planning a trip. For example, an application for booking a flight may be associated with the parameter “end location,” which may be the same parameter of action structure **808** “end location” associated with the rideshare application.

Further, in some examples, the parameter of action structure **808** is a reference to a second application different from the application associated with natural language model **816**. In particular, in addition to the parameter of action structure **808** associated with a first application and also being associated with a second application as discussed above, the parameter of action structure **808** may reference the second application. In some examples, the parameter of action structure **808** references a second parameter of a second action structure associated with the second application.

Continuing the example discussed above where action structure **808** is associated with a rideshare application and the parameter of action structure **808** is “end location” for performing the task of booking a ride share, a second action structure may be associated with a ticket buying application. Accordingly, the parameter “end location” of action structure **808** may reference a second parameter of “location” of a second action structure associated with the ticket buying location. Thus, the parameter “end location” for booking a rideshare may reference the parameter “location” associated with the task of buying movie tickets.

In some examples, the parameter of action structure **808** that references the second parameter of a second action structure associated with a second application may use a value assigned to the second parameter as a value of the parameter of action structure **808**. For example, as discussed above the parameter “end location” for booking a rideshare may reference the parameter “location” associated with the task of buying movie tickets. Accordingly, when a value of “location” is an address of the movie theater at which a user bought tickets, the parameter “end location” may use the address of the movie theater as the value of “end location” for executing a task of booking a ride share.

In some examples, the parameter of action structure **808** references a plurality of parameters of a plurality of action structures associated with a plurality of applications. For example, “end location” of action structure **808** for booking a rideshare may reference any number of “location” parameters associated with various applications including applications for buying movie tickets, making a restaurant reservation, planning a trip, etc. Accordingly, “end location” of

action structure **808** may use the value of any of the plurality of “location” parameters as a value for end location when booking a ride share.

In some examples, natural language model generator **804** automatically determines parameters that the parameter of action structure **808** references based on action structure **808**. For example, when natural language model generator **804** receives the action structure “Rideshare: [end location],” as discussed above, natural language model generator **804** may parse the action structure and recognize that the parameter “end location” is a location and thus will reference other locations and other “location” parameters of other action structures.

In some examples, natural language model generator **804** automatically determines parameters that the parameter of action structure **808** references based on action structure **808** and utterance **802**. For example, when natural language model generator **804** receives the utterance “Get me a ride to SFO,” as discussed above, natural language model generator **804** may parse the utterance and recognize that “SFO” is a value of the “end location” parameter of action structure **808**. Accordingly, natural language model generator **804** can determine that the “end location” parameter being filled by “SFO” will reference other locations and other “location” parameters of other action structures.

In some examples, natural language model generator **804** performs action structure association **806** based on utterance **802** and a default or previously provided action structure **808**. In particular, rather than receiving action structure **808** with utterance **802** and the associations between the two from a user or developer, natural language model generator **804** can parse utterance **802** to determine which action structures may correspond to utterance **802**.

In some examples, associating utterance **802** with action structure **808** includes determining one or more parameters of utterance **802**. For example, when natural language model generator **804** receives utterance **802** “Get me a ride to SFO,” natural language model generator **804** may recognize that “SFO” is an abbreviation for San Francisco International Airport and thus determine that “SFO” is a value of a location parameter. Accordingly, natural language model generator **804** may understand that utterance **802** should be mapped to action structures that include at least a “location” parameter.

In some examples, associating utterance **802** with action structure **808** includes determining whether default action structures match or substantially match utterance **802**. Continuing the example discussed above, natural language model generator **804** may search action structures already included in natural language model generator **804** that include a “location” parameter and thus exclude any default action structures that do not include a “location” parameter. Further, natural language model generator **804** may recognize the use of “ride” in utterance **802** and search for action structures which include the word “ride” or similar words in the template of the action structure. Accordingly, natural language model generator **804** may identify that action structure **808** (e.g., “Rideshare: [end location]”) includes both the word “ride” and a “location” parameter and thus that action structure **808** substantially matches utterance **802**.

In some examples, associating utterance **802** with action structure **808** includes identifying (e.g., replacing) a word of utterance **802** with a parameter of action structure **808**. For example, when natural language model generator **804** parses utterance **802** and recognizes that “SFO” is a location, natural language model generator **804** may replace “SFO” in

utterance **802** with a location parameter of an action structure. In particular, natural language model generator **804** may replace “SFO” with “end location” from action structure **808** to determine whether using “SFO” as an “end location” in the task of calling a rideshare would be valid. Thus, an intermediate utterance of “Get me a ride to [end location]” may be generated as part of associating utterance **802** with action structure **808**. Accordingly, natural language model generator **804** can generate one or more intermediate utterances from utterance **802** that include one or more possible parameters from action structure **808** or other parameters.

Further, in some examples, associating utterance **802** with action structure **808** includes determining parameters of determined action structure **808** that reference parameters of other action structures, as discussed above.

After associating action structure **808** with utterance **802**, natural language model generator **804** performs utterance augmentation **810**, as shown in FIG. **8**. In some examples, utterance augmentation **810** includes determining a plurality of augmented utterances **812** based on utterance **802**. Examples of augmented utterances **812** determined by natural language model generator **804** are provided in Table 1 below.

part of speech as the original word of utterance **802**. For example, because “me” is a pronoun, natural language model generator **804** may replace “me” with other pronouns such as “him,” “her,” “them,” or “us.”

In some examples, the related word is based on one or more words of utterance **802** surrounding the original word to be replaced. For example, natural language model generator **804** may determine that based on the word “SFO” in utterance **802**, utterance **802** is likely related to land or air travel and thus may not change “ride” to “float” because “float” is related to travel over water.

In some examples, natural language model generator **804** replaces the word of utterance **802** with each of the possible related words to create the plurality of augmented utterances **812**. For example, natural language model generator **804** may generate an augmented utterance with each of the pronouns that can replace “me.” Thus, the plurality of augmented utterance **814** may include the individual utterances “get her a ride to SFO,” “get him a ride to SFO,” “get them a ride to SFO,” and “get us a ride to SFO.” This process of replacing words to generate augmented utterances is repeated for each related word to create different augmented utterances of augmented utterances **812**.

Augmented Utterance						
Received Utterance	Change Word	Delete Word	Change Phrase	Change Parameter	Change Value of Parameter	Rearrange Utterance
Get me a ride to SFO	Obtain me a ride to SFO	Get me a ride	Find a car to SFO	Get me a ride from SFO	Get me a ride to the movies	A ride to SFO, please
	Get me a lift to SFO	Get a ride to SFO		Get me from here to SFO	Get me a ride to a restaurant	
	Get her a ride to SFO					
Where is my car?	Where is her car? Where is my ride?	Where's the car?	Is that my car?	Where is SFO?	Where is my limousine?	Is my car here?
Make me a reservation for Chinese food	Make us a reservation for Chinese food	Make me a reservation for Chinese	Get us a time for Chinese food	Make me a reservation for the movies	Make me a reservation for Italian food	Make me a Chinese food reservation
	Get me a reservation for Chinese food	Get me Chinese food		Make me a reservation for a flight	Make me a reservation for Greek food	

In some examples, determining the plurality of augmented utterances **812** based on utterance **802** includes changing a word of utterance **802** to a related word to generate an augmented utterance. In some examples, the related word is a synonym of the original word of utterance **802**. For example, when utterance **802** is “Get me a ride to SFO,” natural language model generator **804** may replace “get” with “obtain” or “ride” with “lift” because “obtain” and “lift” are synonyms of “get” and “ride” respectively. In some examples, the related word is a word with a similar meaning to the original word of utterance **802**. For example, natural language model generator **804** may replace “get” with “order” or “ride” with “car” because “order” and “car” have similar meanings to “get” and “ride” respectively.

In some examples, the related word is an antonym of the original word of utterance **802**. For example, natural language model generator **804** may replace “me” with “you.” In some examples, the related word is a word that is the same

In some examples, natural language model generator **804** replaces multiple words of utterance **802** with related words to generate an augmented utterance. Thus, in addition to the augmented utterances created by replacing a single word of utterance **802** as discussed above, the plurality of augmented utterances **812** also includes augmented utterances with two or more terms different from utterance **802**. For example, natural language model generator **804** may generate an augmented utterance “Order her a car to SFO,” replacing “get” with “order,” “me” with “her,” and “ride” with “car,” to determine the augmented utterance based on utterance **802**. This process is also repeated for different combinations of the different replaced words. Accordingly, many different augmented utterances are determined (e.g., generated) to include different possible combinations of different related words and original words from utterance **802**.

In some examples, determining the plurality of augmented utterances **812** based on utterance **802** includes

deleting a word of utterance **802**. For example, natural language model generator **804** may delete “SFO,” from utterance **802** to create augmented utterance “get me a ride.” Similarly, natural language model generator **804** may delete “me” from utterance **802** to create augmented utterance “get a ride to SFO.” In some examples, deleting a word of utterance **802** includes merging one or more words of utterance **802** together. For example, when the utterance “Where is my car?” is received natural language model generator **804** may delete it by merging “where” and “is” to create the augmented utterance “where’s my car?.”

In some examples, determining the plurality of augmented utterances **812** based on utterance **802** includes changing a phrase of utterance **802** to a related phrase to generate an augmented utterance. For example, natural language model generator **804** may replace “Get me a ride,” of utterance **802** with “Find a car to” to create the augmented utterance “Find a car to SFO.” As another example, natural language model generator **804** may replace “Where is my” of the utterance “where is my car?” with “Is that my” to create the augmented utterance “Is that my car?.” Thus, natural language model generator **804** may replace a phrase of utterance **802** with a related phrase (e.g., a phrase having a similar meaning) to determine a variety of augmented utterances.

In some examples, determining the plurality of augmented utterances **812** based on utterance **802** includes changing a parameter of utterance **802** to a related parameter to generate an augmented utterance. For example, natural language model generator **804** may replace the parameter “end location” with a parameter of “start location” as well as replacing one or more words of utterance **802** to determine an augmented utterance. In particular, natural language model generator **804** may generate the augmented utterance of “Get me a ride from [start location],” from utterance **802** by changing “to” to “from” and changing the parameter of “end location” to “start location.”

In some examples, determining the plurality of augmented utterances **812** based on utterance **802** includes changing a value of the parameter of utterance **802** to another possible value of the parameter to generate an augmented utterance. For example, natural language model generator **804** may replace “SFO” of utterance **802** with other values that could fill an “end location” parameter such as restaurants, movie theaters, other airports, venues, and other land marks. In some examples, the replacement values are a set of default values used when a type of parameter is included in utterance **802**. For example, when a location parameter is included in an utterance, natural language model generator **804** may have a default list of values that could fill a location parameter to use to determine augmented utterances. In some examples, the replacement values are a set of values representing places that are most commonly requested by a set of users.

For example, natural language model generator **804** may change the value of a location parameter to various locations that are commonly requested by users of similar applications or user devices.

In some examples, determining the plurality of augmented utterances **812** based on utterance **802** includes changing the value of the parameter of utterance **802** to a reference to a parameter associated with a second application to generate an augmented utterance. For example, natural language model generator **804** may recognize that “SFO” is a location and change “SFO” to a reference to a location parameter of an application, such as an application for buying movie tickets as discussed above. In some

examples, changing the value of the parameter of utterance **802** to a reference to a parameter associated with a second application is repeated for other parameters associated with third, fourth, fifth, or any number of applications. Accordingly, when a value such as “SFO” is recognized as a location, an augmented utterance replacing “SFO” with references to location parameters of many different applications, allowing natural language model generator **804** to generate a plurality of augmented utterances that are capable of interacting with applications that are associated with parameters similar to the parameters of utterance **802**.

In some examples, determining the plurality of augmented utterances **812** based on utterance **802** includes changing the value of the parameter of utterance **802** to a reference to a parameter associated with the second application determined from a search task associated with the second application. For example, when natural language model generator **804** recognizes that “SFO” is a location and marks “SFO” as a possible reference to an application such as the application for buying movie tickets, natural language model generator **804** may query the movie ticket application. Thus, natural language model generator **804** may query the movie ticket application or a natural language model associated with the movie ticket application for parameters that may satisfy the “location” reference identified by natural language model generator **804** in utterance **802**. The movie ticket application may then return the location of the closest movie theater to satisfy the location parameter of action structure **808** and include the location of the closest movie theater in one or more augmented utterances of the plurality of augmented utterances **812**.

In some examples, determining the plurality of augmented utterances **812** based on utterance **802** includes restructuring the utterance to generate an augmented utterance. In some examples, restructuring the utterance includes changing the order of words in the utterance. For example, natural language model generator **804** may restructure an utterance from “Get me a ride to SFO from work,” to “Get me a ride from work to SFO.”

The plurality of augmented utterances may be determined (e.g., generated) by using one or more of the method discussed above to change utterance **802** into a plurality of augmented utterances. Each of these augmented utterances can then be added to natural language model **816**, as discussed below, to create a complex natural language model capable of recognizing many different user requests. Further, this can be done based on a select set of input utterances **802**, such system **800** can generate natural language model **816** based on a relatively small sample of inputs. This results in a system that is more responsive to user inputs because natural language model **816** includes augmented utterances which may not typically be included in a natural language model created using conventional methods.

After determining the plurality of augmented utterances **812** based on utterance **802**, natural language model generator **804** performs utterance mapping **814**, as shown in FIG. 8.

In some examples, performing utterance mapping **814** includes creating natural language model **816** including the received utterance **802** and the plurality of augmented utterances **812** by mapping the plurality of augmented utterances **812** to the associated action structure **808**. In some examples, mapping the plurality of augmented utterances **812** to the associated action structure **808** includes mapping each of the augmented utterances **812** to action structure **808**. For example, natural language model generator **804** may map each of the augmented utterances “Order

her a car to SFO,” “get us a ride to SFO” and “Find me a ride to the movies,” to action structure **808** of “Rideshare: [end location].” Accordingly, each of the augmented utterances is mapped to the action structure allowing natural language model **814** to determine the appropriate intent or task if any of the augmented utterances are received.

In some examples, mapping each of the augmented utterances **812** to action structure **808** includes mapping words of each of the augmented utterances **812** to a template of action structure **808**. For example, natural language model generator **804** may map the augmented utterance “order her a car to the movies,” to action structure **808** of “Rideshare: [end location]” by mapping “order her a car to” to the task of “rideshare” and “the movies” to the parameter of “end location.” Similarly, natural language model generator **804** may map the augmented utterance “obtain me a lift to SFO,” to action structure **808** by mapping “obtain me a lift to the task of “rideshare” and “SFO” to the parameter of “end location.”

In some examples, mapping words of each of the augmented utterances **812** to a template of action structure **808** includes mapping an augmented parameter to a parameter of action structure **808**. For example, as discussed above, natural language model generator **804** may map “the movies” of augmented utterance “order her a car to the movies,” to the parameter “end location” of action structure **808**. In some examples, mapping words of each of the augmented utterances **812** to a template of action structure **808** includes mapping a reference to a parameter of another application to a parameter of action structure **808**. For example, natural language model generator **804** may map the augmented utterance “get me a ride to [location],” where [location] is a reference to a parameter of an action structure associated with an application for making restaurant reservations by mapping “location” (e.g., the reference to a parameter of another application) to the parameter “end location” of action structure **808**.

In some examples, mapping an augmented parameter to a parameter of action structure **808** includes inserting the augmented parameter into action structure **808**. In some examples, inserting the augmented parameter into action structure **808** includes determining whether the task of action structure **808** could be completed with the augmented parameter. For example, when natural language model generator **804** maps “the movies” of augmented utterance “order her a car to the movies,” to the parameter “end location” of action structure **808**, natural language model generator **804** can insert “the movies” into action structure **808** and determine if the task of “rideshare” can be completed.

In some examples, mapping a reference to a parameter of another application to a parameter of action structure **808** includes inserting the reference to the parameter of another application into action structure **808**. In some examples, inserting the reference to the parameter of another application into action structure **808** includes determining whether the task of action structure **808** could be completed with the reference to the parameter of another application. For example, when natural language model generator **804** maps the reference to a parameter of an application for making restaurant reservations of “location” to action structure **808**, natural language model generator **804** can insert the reference “location” into action structure **808** and determine if the task of “rideshare” can be completed.

Accordingly, by mapping each of the augmented utterances **812** to action structure **808** natural language model generator **804** creates associations between each of the augmented utterance **814** and action structure **808** that

mirror the association between utterance **802** and action structure **808**. These associations, utterance **802**, augmented utterance **814**, and action structure **808** are all combined to create natural language model **816**.

In some examples, natural language model generator **804** is a neural network, machine learning model, or simple network used to process data. Accordingly, in some examples, the various processes of natural language model generator **804** discussed above are implemented as functions of a neural network or machine learning model.

After performing utterance mapping **814**, natural language model generator **804** provides natural language model **816** to an electronic device. In some examples, providing natural language model **816** to an electronic device includes integrating natural language model **816** into a plurality of natural language models. Further, as discussed above, natural language model **816** is associated with a particular application. Accordingly, each of the plurality of natural language models are associated with a different application. For example, as discussed above, natural language model **816** may be associated with a ride share application. Therefore, a natural language model of the plurality of natural language models may be associated with a restaurant reservation application, another natural language model of the plurality of natural language models may be associated with a movie ticket buying application, etc.

Thus, natural language model **816** may be integrated with the plurality of natural language models to provide a combined set of natural language models that cover a plurality of applications. In some examples, the plurality of natural language models corresponds to a plurality of applications installed on the electronic device. For example, natural language model **816** may be provided to the electronic device when the rideshare application associated with natural language model **816** is installed on the electronic device. Accordingly, natural language model **816** may be integrated with the natural language models for the other applications already installed on the electronic device.

In some examples, providing natural language model **816** to an electronic device includes integrating natural language model **816** with a natural language model of a digital assistant. For example, when natural language model **816** is provided to an electronic device, the electronic device may include a digital assistant. Accordingly, the digital assistant may integrate natural language model **816** with the digital assistant’s own native natural language understanding capabilities to allow a user to interact with the application associated with natural language model **816** through the digital assistant.

In some examples, providing natural language model **816** to an electronic device includes providing natural language model **816** to a user device. For example, the digital assistant discussed above may be implemented on a user device such as a phone, watch, tablet, etc. Accordingly, natural language model **816** may be provided to that user device to be integrated with the digital assistant (or other natural language models) so that a user may access and utilize natural language model **816** on the user device.

As discussed above, in some examples, natural language model generator **804** receives a second utterance and processes the second utterance in the same manner as the first utterance to create natural language model **816**. Thus, natural language model generator **804** performs action structure association **806**, utterance augmentation **810**, and utterance mapping **814** on the second utterances as discussed above with regard to utterance **802** to create natural language model **816**. In particular, natural language model generator

804 associates a second action structure of the application associated with natural language model **816** with the second utterance and then determines a plurality of augmented utterances based on the second utterance. Natural language model generator **804** then maps the plurality of augmented utterance based on the second utterance to the second action structure and updates natural language model **816** with the second utterance and the plurality of augmented utterances based on the second utterance. The updated natural language model **816** may then be provided to another electronic device, digital assistant, or integrated with other natural language models.

For example, as discussed above, a second utterance of “where is my car?” may be received by natural language model generator **804**. This utterance may then be associated to an action structure for requesting information about a ride share of “Information: [ride share],” by associating “where is” with the task of determining information and “my car” with the parameter of “ride share.” The utterance may then be augmented by changing words, phrases, rearranging, etc. as discussed above including determining an augmented utterance of “where is my ride?.” The augmented utterance may then be mapped to the action structure of “Information: [ride share],” by mapping “where is” to the task of determining information and “my ride” to the parameter of “ride share.” This augmented utterance, as well as other determined and mapped augmented utterances, may then be added to natural language model **816** before it is provided to an electronic device.

In some examples, natural language model generator **804** processes the second utterance after completing processing on the first utterance. Therefore, natural language model **816** may be determined iteratively by processing multiple utterances in succession. In some examples, natural language model generator **804** process the second utterance while simultaneously processing the first utterance. Thus, natural language model **816** can be determined by processing multiple utterances at the same time when they are provided together.

It will be understood that this process can be repeated for any number of utterances that a user or developer wishes to be augmented and processed for inclusion in natural language model **816**. Thus, system **800** could receive and process three, four, five, six, or any other feasible number of utterances to create a plurality of augmented utterances to be included in natural language model **816** either iteratively or simultaneously.

In some examples, after creating and providing natural language model **816**, natural language model generator **804** receives a third utterance to process for creating a second natural language model. Thus, natural language model generator **804** processes the third utterance to create a second natural language model by performing action structure association **806**, utterance augmentation **810**, and utterance mapping **814**. In particular, natural language model generator **804** associates a third action structure of a second application associated with the second natural language model with the third utterance and then determines a plurality of augmented utterances based on the third utterance. Natural language model generator **804** then maps the plurality of augmented utterance based on the third utterance to the third action structure and creates the second natural language model with the third utterance and the plurality of augmented utterances based on the third utterance. The second natural language model may then be provided to another electronic device, digital assistant, or integrated with other natural language models. In some examples, the sec-

ond natural language model is provided with natural language model **816**. In some examples, the second natural language model is provided separate from natural language model **816**.

For example, natural language model generator **804** may receive the utterance “make me a reservation for Chinese food.” Accordingly, natural language model generator **804** may associate the utterance with an action structure for a restaurant reservation application of “Reservation: [type of restaurant],” because of the term “reservation” included in the utterance and because “Chinese food” is a type or category of restaurant. Natural language model generator **804** may then augment the utterance to create augmented utterances like “make me a reservation for Italian food,” and “make me a reservation for Greek food.” These augmented utterances may then be mapped to the action structure of “Reservation: [type of restaurant],” by mapping each of the augmented parameters to “type of restaurant.” These augmented utterances and the received utterance are then added to a natural language model for the restaurant reservation application which can be provided to another electronic device, digital assistant, or integrated with other natural language models.

As discussed above with regard to multiple utterances for the first application, this process may be completed with multiple utterances for the second application as required to create the desired natural language model. Additionally, this process may be repeated multiple times for multiple utterances of different applications to create a plurality of natural language models where each natural language model is associated with a different application.

Accordingly, an ecosystem of complex and interconnected natural language models may be determined automatically based on a relatively small sample of utterances, increasing the responsiveness and efficiency of a natural language system or digital assistant with less work by a developer or user creating the natural language models.

FIG. 9A-9B illustrate process **900** for creating or augmenting a natural language model, according to various examples. Process **900** is performed, for example, using one or more electronic devices including system **800**. In some examples, process **900** is performed using a client-server system (e.g., system **100**), and the blocks of process **900** are divided up in any manner between the server (e.g., DA server **106**) and a client device. In other examples, the blocks of process **900** are divided up between the server and multiple client devices (e.g., a mobile phone and a smart watch). Thus, while portions of process **900** are described herein as being performed by particular devices of a client-server system, it will be appreciated that process **900** is not so limited. In other examples, process **900** is performed using only a client device (e.g., user device **104**) or only multiple client devices. In process **900**, some blocks are, optionally, combined, the order of some blocks is, optionally, changed, and some blocks are, optionally, omitted. In some examples, additional steps may be performed in combination with the process **900**.

At block **902**, an utterance (e.g., utterance **802**) is received. In some examples, the utterance is part of a training data for a natural language model (e.g., natural language model **816**) of an application.

At block **904**, an action structure (e.g., action structure **808**) is associated with the utterance (e.g., utterance **802**). In some examples, the action structure comprises a template including one or more parameters. In some examples, the parameter is associated with the application. In some examples, the parameter is a reference to a second applica-

tion. In some examples, the reference is a second parameter associated with the second application.

At block **906**, a plurality of augmented utterances (e.g., plurality of augmented utterances **812**) are determined based on the received utterance (e.g., utterance **802**). At block **908**, determining the plurality of augmented utterances based on the received utterance further comprises changing a word of the utterance to a related word. At block **910**, determining the plurality of augmented utterances based on the received utterance further comprises changing a phrase of the utterance to a related phrase. At block **912**, determining the plurality of augmented utterances based on the received utterance further comprises changing a parameter of the utterance to a related parameter

At block **914**, determining the plurality of augmented utterances (e.g., plurality of augmented utterances **812**) based on the received utterance (e.g., utterance **802**) further comprises changing a value of the parameter to another possible value of the parameter. At block **916**, changing the value of the parameter further comprises changing the parameter to a reference to a natural language model associated with another application. At block **918**, determining the plurality of augmented utterances based on the received utterance further comprises restructuring the utterance.

At block **920**, a natural language model (e.g., natural language model **816**) including the received utterance (e.g., utterance **802**) and the plurality of augmented utterances (e.g., plurality of augmented utterances **812**) is created by mapping the plurality of augmented utterance to the associated action structure (e.g., action structure **808**). At block **922**, mapping the plurality of augmented utterances to the associated action structure further comprises mapping an augmented parameter to a parameter of the action structure. At block **924**, mapping the plurality of augmented utterances to the associated action structure further comprises mapping a reference to a second application to the associated action structure. At block **926**, mapping the reference to the second application into the associated action structure further comprises mapping the reference to a parameter of the action structure.

At block **928**, the natural language model (e.g., natural language model **816**) including the received utterance (e.g., utterance **802**) and the plurality of augmented utterances (e.g., plurality of augmented utterances **812**) is provided to a second electronic device. At block **930**, providing the natural language model to the second electronic device further comprises integrating the natural language model into a plurality of natural language models.

In some examples, a second utterance (e.g., utterance **802**) is received and a second action structure (e.g., action structure **808**) of the application is associated with the second utterance. In some examples, a second plurality of augmented utterances (e.g., plurality of augmented utterances **812**) is determined based on the second utterance. In some examples, the natural language model (e.g., natural language model **816**) is updated with the second utterance and the second plurality of augmented utterances by mapping the second plurality of augmented utterances to the second action structure. In some examples, the updated natural language model is provided to the second electronic device.

In some examples, a third utterance (e.g., utterance **802**) is received and a third action structure (e.g., action structure **808**) of a third application is associated with the third utterance. In some examples, a third plurality of augmented utterances (e.g., plurality of augmented utterances **812**) is determined based on the third utterance. In some examples, a second natural language model (e.g., natural language

model **816**) is created with the third utterance and the third plurality of augmented utterances by mapping the third plurality of augmented utterances to the third action structure. In some examples, the second natural language model is provided to the second electronic device.

The operations described above with reference to FIGS. **9A-9B** are optionally implemented by components depicted in FIGS. **1-4**, **6A-B**, **7A-C**, and **8**. It would be clear to a person having ordinary skill in the art how other processes are implemented based on the components depicted in FIGS. **1-4**, **6A-B**, **7A-C**, and **800**.

FIG. **10** illustrates a process **1000** for updating a natural language model, according to various examples. Process **1000** is performed, for example, using one or more electronic devices implementing a digital assistant. In some examples, process **1000** is performed using a client-server system (e.g., system **100**), and the blocks of process **1000** are divided up in any manner between the server (e.g., DA server **106**) and a client device. In other examples, the blocks of process **1000** are divided up between the server and multiple client devices (e.g., a mobile phone and a smart watch). Thus, while portions of process **1000** are described herein as being performed by particular devices of a client-server system, it will be appreciated that process **1000** is not so limited. In other examples, process **1000** is performed using only a client device (e.g., user device **104**) or only multiple client devices. In process **1000**, some blocks are, optionally, combined, the order of some blocks is, optionally, changed, and some blocks are, optionally, omitted. In some examples, additional steps may be performed in combination with the process **1000**.

Generally, operations of process **1000** may be implemented to update a natural language model based on an utterance including a word unrecognized by the natural language model. As described in further detail below, an utterance may be received by an electronic device, and in particular by a digital assistant of an electronic device. One or more words of the utterance unrecognized by a first natural language model associated with a first application may then be determined and marked as a reference. Whether a parameter of a second natural language model associated with a second application satisfies a requirement of the reference can then be determined. If the parameter of the second natural language model associated with the second application satisfies the requirement of the reference, the relationship between the reference and the parameter in the first natural language model associated with the first application may be stored.

At block **1010**, an utterance is received (e.g., at electronic device **700**). In some examples, the utterance is received from a user. In some examples, the utterance is directed to a digital assistant (e.g., the digital assistant of electronic device **700**). In some examples, the utterance is a request related to an application. For example, the utterance may be the request "Assistant, get me a ride to my next meeting," that is directed to a digital assistant and is related to the application of getting a ride share. In some examples, the application is determined by determining a task of the utterance. For example, when the utterance includes "get me a ride" the digital assistant and/or the natural language model may determine that a task of the utterance is to get the user a ride share to a location.

In some examples, the utterance is received from another device, such as a device communicatively coupled to the electronic device (e.g., smart watch). In some examples, the user input is provided over an ad-hoc connection between the electronic device and the other device. In other

examples, the user input is provided over a multi-hop network, such as the Internet.

At block **1020**, one or more words of the utterance that are unrecognized by a first natural language model associated with a first application are determined. In some examples, determining the one or more words of the utterance that are unrecognized by the first natural language model includes providing the utterance to the first natural language model. In some examples, after receiving the utterance, the first natural language model processes the utterance to determine a task of the utterance. For example, as discussed above, once the utterance “get me a ride to my next meeting,” is received it may be provided to the first natural language model which may process the utterance and determine based on the user of “get me a ride” that a task of the utterance is to get a ride share.

In some examples, after receiving the utterance, the first natural language model processes the utterance to determine whether each of the words of the utterance is recognized by the natural language model and/or related to the first application. In some examples, determining the one or more words of the utterance that are unrecognized by the first natural language model includes mapping the utterance to an action structure associated with the first natural language model. Continuing the previous example, the first natural language model may map “get me a ride” to the task “rideshare” of an action structure “Rideshare: [end location].” Further, the first natural language model may not understand the use of “to my next meeting” and therefore may not be able to map it to the parameter “end location” of the action structure. Accordingly, the first natural language model may determine that “to my next meeting” are not recognized (e.g., in the context of the selected action structure).

In some examples, determining the one or more words of the utterance that are unrecognized by the first natural language model includes selecting the first natural language model based on one or more words of the utterance. In some examples, selecting the first natural language model based on one or more words of the utterance includes processing the utterance with another natural language model associated with the digital assistant to determine whether the utterance is related to the first application. For example, the digital assistant may have a natural language model dedicated to determining a general task of the utterance without doing any detailed processing. Accordingly, the digital assistant may use its natural language model to determine that the utterance of “get me a ride to my next meeting,” is associated with a ride share based on the use of “ride.” Thus, digital assistant may select the first natural language model and provide the utterance to the first natural language model for further processing.

In some examples, determining the one or more words of the utterance that are unrecognized by the first natural language model includes providing the utterance to a plurality of natural language models associated with applications installed on the device. In some examples, providing the utterance to the plurality of natural language models includes performing preliminary processing with each of the plurality of the natural language models to determine whether the utterance is associated with each of the plurality of the natural language models. For example, in addition to the natural language models discussed above, the utterance may be provided to natural language models associated with movie tickets buying applications, restaurant reservation making applications, etc. Accordingly, these applications may parse the utterance and quickly determine that the

utterance is not related to them because the words of the utterance are not related to the tasks of buying movie tickets or making a restaurant reservation.

In some examples, a natural language model of the plurality of natural language models is selected based on the amount of the utterance that is understood by the natural language model. In some examples, the natural language model of the plurality of natural language models is selected based on the amount of the utterance that can be mapped correctly to an action structure of the natural language model. For example, when the utterance “find Chinese food near me” is received a first natural language model associated with a restaurant reservation application may map “find Chinese food” to the action structure of “find restaurant: [food type], [location],” because of the use of “food” and the use of “Chinese” in the utterance and mark “near me” as unrecognized words. A second natural language model associated with translation may also attempt to map the utterance to an action structure based on the use of “Chinese” in the utterance. However, the second natural language model will mark “food” as an unrecognized word in addition to “near me.” Thus, the first natural language model may be selected because the first natural language model recognizes more of the words of the utterance and is able to map more of the words of the utterance to an action structure of the first natural language model.

In some examples, the first natural language model was previously created by augmenting one or more utterances included in training data for the first application, as described above. In some examples, the first natural language model is created using system **800** and natural language model **804** discussed above. In particular, the first natural language model is generated by performing action structure association **806**, utterance augmentation **810**, and utterance mapping **814** on one or more utterances provided as training data for a natural language model associated with a particular application. In some examples, the first natural language model is received from another electronic device and integrated with a plurality of natural language models prior to receiving the utterance. In some examples, the first natural language model is received from another electronic device and integrated with a digital assistant prior to receiving the utterance.

In some examples, augmenting one or more utterances included in training data for the first application includes changing one or more words of the utterances to a related word. In some examples, augmenting one or more utterances included in training data for the first application includes changing one or more phrases of the utterances to a related phrase. In some examples, augmenting one or more utterances included in training data for the first application includes changing a parameter of the utterances to a related parameter. In some examples, augmenting one or more utterances included in training data for the first application includes changing a value of the parameter of the utterances to another possible value of the parameter. In some examples, augmenting one or more utterances included in training data for the first application includes changing the parameter to a reference to another natural language model associated with another application. In some examples, augmenting one or more utterances included in training data for the first application includes restructure the one or more utterances.

At block **1030**, the one or more unrecognized words is marked as a reference (e.g., a reference to another application). In some examples, marking the one or more unrecognized words as a reference includes marking the one or

more unrecognized words as a parameter. In some examples, the parameter is a parameter of a template of an action structure of the natural language model. For example, when the first natural language model processes “get me a ride to my next meeting” and marks “to my next meeting” as unrecognized words, the first natural language model may further determine that “to my next meeting” is a possible parameter value for “end location” of the action structure “rideshare: [end location]” and thus mark “to my next meeting” as the parameter “end location.” In this way, the first natural language model may flag words that could possibly include parameters for executing identified tasks, even if the first natural language model does not understand the specific words marked as a parameter.

After recognizing that the one or more unrecognized words are a possible parameter for completing the task, the first natural language model determines that another natural language model (e.g., the second natural language model) can include a value for the unrecognized words marked as a parameter. Accordingly, the first natural language model marks the unrecognized words as a reference to another application (e.g., a reference to a natural language model associated with another application) and then proceeds to determine which parameter of the other natural language models available may be associated with the unrecognized words, as described below.

At block 1040, whether a parameter of a second natural language model associated with a second application satisfies a requirement of the reference is determined. As a part of this process the first natural language model may interact with the second natural language model or the digital assistant to determine possible parameters of the second natural language model (and other natural language models) could be mapped to the marked unrecognized words. In this way, the first natural language model, with the help of the digital assistant and the other natural language models, establishes a connection to determine a meaning of the unrecognized words.

In some examples, the second natural language model was previously created by augmenting one or more utterances included in training data for the second application. In some examples, the second natural language model is created using system 800 and natural language model 804 discussed above. In particular, the second natural language model is generated by performing action structure association 806, utterance augmentation 810, and utterance mapping 814 on one or more utterances provided as training data for a natural language model associated with a particular application. In some examples, the second natural language model is received from another electronic device and integrated with a plurality of natural language models prior to receiving the utterance. In some examples, the second natural language model is received from another electronic device and integrated with a digital assistant prior to receiving the utterance.

In some examples, augmenting one or more utterances included in training data for the second application includes changing one or more words of the utterances to a related word. In some examples, augmenting one or more utterances included in training data for the second application includes changing one or more phrases of the utterances to a related phrase. In some examples, augmenting one or more utterances included in training data for the second application includes changing a parameter of the utterances to a related parameter. In some examples, augmenting one or more utterances included in training data for the second application includes changing a value of the parameter of the

utterances to another possible value of the parameter. In some examples, augmenting one or more utterances included in training data for the second application includes changing the parameter to a reference to another natural language model associated with another application. In some examples, augmenting one or more utterances included in training data for the second application includes restructure the one or more utterances.

In some examples, determining whether a parameter of a second natural language model associated with a second application satisfies a requirement of the reference includes determining whether the parameter matches one or more properties of the reference. For example, when the first natural language model marks “to my next meeting” as the possible parameter “end location” and a reference to other applications, the reference is given properties such as the parameter type location, possible parameter values including addresses, coordinates, etc, and the application types of applications for planning trips, reserving restaurants, and other applications that would have some association to possible end locations for a rideshare task.

In some examples, the property of the reference is a parameter type. Exemplary parameter types include, but are not limited to, location, color, orientation, food type, time, month, day, year, temperature, car type, or any other possible type of parameter. In some examples, the property of the reference is a parameter value. The parameter value includes any valid value that could be provided to the parameter for a task to be completed. For example, for the parameter type of location a parameter value could be an address or a set of coordinates. As another example, for the parameter type of orientation a parameter value could be north, south, up, down, an angle, etc. In some examples, the property of the reference is an application type. Exemplary application types include categories of application such as flight booking applications, restaurant applications, word processing applications, gaming applications, scheduling applications, messaging applications, etc.

Continuing the example above, when “to my next meeting” is marked as a reference and the parameter “end location,” the first natural language model may determine, with the help of the digital assistant and/or the second natural language model, that the second natural language model and the application associated with the second natural language model may be able to provide a value for the parameter. In particular, the digital assistant and/or the second natural language model may detect the use of “meeting” in the reference and determine that the reference is related to the second application which is a scheduling or calendar application.

Accordingly, the second natural language model may provide the parameter “meeting location” associated with the task of scheduling meetings from the second natural language model to the first natural language model. Further, the second natural language model may provide a value for the parameter of “meeting location” including an address where the user is supposed to have their next meeting. The first natural language model may then determine that the parameter “meeting location” satisfies the requirements of the reference because a location parameter was required, or because the address is a possible value for the reference, or because the scheduling application is one of the applications that meets the requirements of the reference.

In some examples, the second natural language model or another natural language model may determine a possible parameter by processing a query associated with the second natural language model that may satisfy the reference and/or

61

one or more requirements of the reference. For example, the second natural language model associated with a calendar may perform a search query to determine a value that will satisfy a location typed parameter, as determined when “my next meeting” is marked as a reference by the first natural language model. Accordingly, a calendar search task (e.g., the query) associated with the second natural language model can determine a location (e.g., an address) associated with “my next meeting,” and provide the address to the first natural language model. This address may then be injected (e.g., mapped) into the reference to determine whether the address satisfies the requirements of the reference as discussed above.

At block 1050, in accordance with a determination that the parameter of the second natural language model associated with the second application satisfies the requirement of the reference, the relationship between the reference and the parameter in the first natural language model associated with the first application is stored.

In some examples, storing the relationship between the reference and the parameter in the first natural language model associated with the first application includes updating the first natural language model associated with the first application to recognize the one or more unrecognized words as a reference to the parameter of the second natural language model. In some examples, updating the first natural language model associated with the first application to recognize the one or more unrecognized words as a reference to the parameter of the second natural language model includes mapping the reference to the parameter of the second natural language model to the associated action structure of the first natural language model. For example, once the parameter “meeting location” of the second natural language model is identified, the first natural language model maps the parameter to the action structure “rideshare: [end location]” by inserting “meeting location” into the parameter “end location.” The first natural language model may then save this mapping as the association so that when the phrase “to my next meeting” is received again, the first natural language model recognizes that it should reference the “meeting location” parameter of the second natural language model.

Further, in some examples, this process may be repeated for multiple other natural language models. In particular, the first natural language model determines whether a parameter of a third natural language model associated with a third application satisfies a requirement of the reference and in accordance with a determination that the parameter of the third natural language model associated with the third application satisfies the requirement of the reference, the relationship between the reference and the parameter in the first natural language model associated with the first application is stored. This may be repeated for a fourth natural language model, a fifth natural language model, etc. until all of the natural language models have been checked or a parameter satisfying the requirement of the reference is determined.

Furthermore, the utterance may also be processed iteratively or simultaneously by multiple applications to determine unrecognized words and determine if the application associated with the respective natural language model should respond or process the utterance. In particular, whether the one or more words of the utterance are unrecognized by a fourth natural language model associated with a fourth application can be determined. Further, the one or more unrecognized words are marked as a second reference associated with the fourth application and whether a parameter of the second natural language model (or another natural

62

language model) satisfies a requirement of the second reference is determined. Then, in accordance with a determination that the parameter of the second natural language model associated with the second application satisfies the requirement of the second reference, the relationship between the second reference and the parameter in the fourth natural language model associated with the fourth application is stored.

Accordingly, it will be understood that the process may be repeated for any number of utterances and any number of natural language models processing those utterances to determine which of the natural language models can process the utterances and can make the proper references to other natural language models. In this way the efficiency of the natural language models, the electronic device, and the digital assistant may be improved when processing natural language inputs. This in turn reduces the amount of processing required and improves the user experience.

In some examples, the relationship between the reference and the parameter in the first natural language model associated with the first application is added to a database of relationships that can be accessed by the various natural language models and the digital assistant to resolve ambiguities and determine the meaning of words that are unrecognized. This database may be accessed whenever an unrecognized word is detected to determine if the unrecognized word has been linked to any of the natural language models previously.

In some examples, after updating the first natural language model, a second utterance is received. Accordingly, the first natural language model determines whether the second utterance includes the reference. For example, the first natural language model may receive the utterance “get me to my next meeting,” after being updated and determine that the utterance includes the phrase “to my next meeting.” Accordingly, the first natural language model can recognize that this is the same phrase previously marked as a reference.

In accordance with a determination that the second utterance includes the reference, the first natural language model replaces the reference with a value of the parameter of the second natural language model. For example, as discussed above, because the association has been saved, the first natural language model understands that “to my next meeting” is referencing the “meeting location” parameter of the second natural language model and replaces this phrase with the value of the “meeting location” parameter from the second natural language model for executing the ride share task.

However, in accordance with a determination that the second utterance does not include the reference, the digital assistant determines one or more words of the second utterance unrecognized by the first natural language model associated with the first application. Accordingly, the digital assistant marks the one or more unrecognized words as a second reference and then may process the one or more unrecognized words as discussed above.

Processing the one or more unrecognized words and determining one or more words of the second utterance unrecognized by the first natural language model associated with the first application and marking the unrecognized words as the second reference is then performed in substantially the same manner as discussed above with regard to determining unrecognized words of the first utterance.

It will be understood that this process may be repeated for any number of unrecognized words or phrases as needed. In this way, the natural language models may be continually updated and connected together, providing a more respon-

sive experience to the user and more efficient processing of natural language inputs by leveraging the assets of the various natural language models together.

In accordance with some implementations, a computer-readable storage medium (e.g., a non-transitory computer readable storage medium) is provided, the computer-readable storage medium storing one or more programs for execution by one or more processors of an electronic device, the one or more programs including instructions for performing any of the methods or processes described herein.

In accordance with some implementations, an electronic device (e.g., a portable electronic device) is provided that comprises means for performing any of the methods or processes described herein.

In accordance with some implementations, an electronic device (e.g., a portable electronic device) is provided that comprises a processing unit configured to perform any of the methods or processes described herein.

In accordance with some implementations, an electronic device (e.g., a portable electronic device) is provided that comprises one or more processors and memory storing one or more programs for execution by the one or more processors, the one or more programs including instructions for performing any of the methods or processes described herein.

The foregoing description, for purpose of explanation, has been described with reference to specific embodiments. However, the illustrative discussions above are not intended to be exhaustive or to limit the invention to the precise forms disclosed. Many modifications and variations are possible in view of the above teachings. The embodiments were chosen and described in order to best explain the principles of the techniques and their practical applications. Others skilled in the art are thereby enabled to best utilize the techniques and various embodiments with various modifications as are suited to the particular use contemplated.

Although the disclosure and examples have been fully described with reference to the accompanying drawings, it is to be noted that various changes and modifications will become apparent to those skilled in the art. Such changes and modifications are to be understood as being included within the scope of the disclosure and examples as defined by the claims.

As described above, one aspect of the present technology is the gathering and use of data available from various sources to improve the processing of user requests with natural language models. The present disclosure contemplates that in some instances, this gathered data may include personal information data that uniquely identifies or can be used to contact or locate a specific person. Such personal information data can include demographic data, location-based data, telephone numbers, email addresses, twitter IDs, home addresses, data or records relating to a user's health or level of fitness (e.g., vital signs measurements, medication information, exercise information), date of birth, or any other identifying or personal information.

The present disclosure recognizes that the use of such personal information data, in the present technology, can be used to the benefit of users. For example, the personal information data can be used to provide personalized tasks to the user. Accordingly, use of such personal information data enables users to calculated control of the provided tasks. Further, other uses for personal information data that benefit the user are also contemplated by the present disclosure. For instance, health and fitness data may be used to

provide insights into a user's general wellness, or may be used as positive feedback to individuals using technology to pursue wellness goals.

The present disclosure contemplates that the entities responsible for the collection, analysis, disclosure, transfer, storage, or other use of such personal information data will comply with well-established privacy policies and/or privacy practices. In particular, such entities should implement and consistently use privacy policies and practices that are generally recognized as meeting or exceeding industry or governmental requirements for maintaining personal information data private and secure. Such policies should be easily accessible by users, and should be updated as the collection and/or use of data changes. Personal information from users should be collected for legitimate and reasonable uses of the entity and not shared or sold outside of those legitimate uses. Further, such collection/sharing should occur after receiving the informed consent of the users. Additionally, such entities should consider taking any needed steps for safeguarding and securing access to such personal information data and ensuring that others with access to the personal information data adhere to their privacy policies and procedures. Further, such entities can subject themselves to evaluation by third parties to certify their adherence to widely accepted privacy policies and practices. In addition, policies and practices should be adapted for the particular types of personal information data being collected and/or accessed and adapted to applicable laws and standards, including jurisdiction-specific considerations. For instance, in the US, collection of or access to certain health data may be governed by federal and/or state laws, such as the Health Insurance Portability and Accountability Act (HIPAA); whereas health data in other countries may be subject to other regulations and policies and should be handled accordingly. Hence different privacy practices should be maintained for different personal data types in each country.

Despite the foregoing, the present disclosure also contemplates embodiments in which users selectively block the use of, or access to, personal information data. That is, the present disclosure contemplates that hardware and/or software elements can be provided to prevent or block access to such personal information data. For example, in the case of personalized tasks and natural language models, the present technology can be configured to allow users to select to "opt in" or "opt out" of participation in the collection of personal information data during registration for services or anytime thereafter. In another example, users can select not to provide location or other personal data for completion of tasks. In yet another example, users can select to limit the length of time location and other personalized data is maintained. In addition to providing "opt in" and "opt out" options, the present disclosure contemplates providing notifications relating to the access or use of personal information. For instance, a user may be notified upon downloading an app that their personal information data will be accessed and then reminded again just before personal information data is accessed by the app.

Moreover, it is the intent of the present disclosure that personal information data should be managed and handled in a way to minimize risks of unintentional or unauthorized access or use. Risk can be minimized by limiting the collection of data and deleting data once it is no longer needed. In addition, and when applicable, including in certain health related applications, data de-identification can be used to protect a user's privacy. De-identification may be facilitated, when appropriate, by removing specific identi-

fiers (e.g., date of birth, etc.), controlling the amount or specificity of data stored (e.g., collecting location data at a city level rather than at an address level), controlling how data is stored (e.g., aggregating data across users), and/or other methods.

Therefore, although the present disclosure broadly covers use of personal information data to implement one or more various disclosed embodiments, the present disclosure also contemplates that the various embodiments can also be implemented without the need for accessing such personal information data. That is, the various embodiments of the present technology are not rendered inoperable due to the lack of all or a portion of such personal information data. For example, natural language models can be created by inferring preferences based on non-personal information data or a bare minimum amount of personal information, such as the content being requested by the device associated with a user, other non-personal information available to the natural language models, or publicly available information.

What is claimed is:

1. An electronic device comprising:
 - one or more processors;
 - a memory; and
 - one or more programs, wherein the one or more programs are stored in the memory and configured to be executed by the one or more processors, the one or more programs including instructions for:
 - receiving an utterance;
 - determining one or more words of the utterance unrecognized by a first natural language model associated with a first application;
 - marking the one or more unrecognized words as a reference;
 - determining whether a parameter of a second natural language model associated with a second application satisfies a requirement of the reference;
 - in accordance with a determination that the parameter of the second natural language model associated with the second application satisfies the requirement of the reference, storing the relationship between the reference and the parameter in the first natural language model associated with the first application.
2. The electronic device of claim 1, wherein the utterance is received from a user.
3. The electronic device of claim 1, wherein the first natural language model was previously created by augmenting one or more utterances included in training data for the first application.
4. The electronic device of claim 1, wherein the second natural language model was previously created by augmenting one or more utterances included in training data for the second application.
5. The electronic device of claim 1, wherein determining whether the parameter of the second natural language model associated with the second application satisfies the requirement of the reference further comprises:
 - determining whether the parameter matches one or more properties of the reference.
6. The electronic device of claim 5, wherein the one or more properties of the reference includes at least one of a parameter type, a parameter value, and an application type.
7. The electronic device of claim 1, wherein storing the relationship between the reference and the parameter in the first natural language model associated with the first application further comprises:
 - updating the first natural language model associated with the first application to recognize the one or more

unrecognized words as a reference to the parameter of the second natural language model.

8. The electronic device of claim 3, wherein augmenting one or more utterances included in training data for the first application further comprises changing one or more words of each of one of the utterances to a related word.

9. The electronic device of claim 1, wherein determining one or more words of the utterance unrecognized by a first natural language model associated with a first application further comprises:

- mapping the one or more words of the utterance to an action structure associated with the first natural language model.

10. The electronic device of claim 9, wherein the action structure associated with the first natural language model comprises a template including one or more parameters.

11. The electronic device of claim 1, wherein the utterance is a first utterance and the reference is a first reference, the one or more programs further including instructions for:

- receiving a second utterance;
- determining whether the second utterance includes the first reference;
- in accordance with a determination that the second utterance includes the first reference, replacing the first reference with the parameter of the second natural language model; and
- in accordance with a determination that the second utterance does not include the reference:
 - determining one or more words of the second utterance unrecognized by a first natural language model associated with a first application; and
 - marking the one or more unrecognized words as a second reference.

12. The electronic device of claim 1, the one or more programs further including instructions for:

- determining whether a parameter of a third natural language model associated with a third application satisfies a requirement of the reference;
- in accordance with a determination that the parameter of the third natural language model associated with the third application satisfies the requirement of the reference, storing the relationship between the reference and the parameter in the first natural language model associated with the first application.

13. The electronic device of claim 1, wherein the reference is a first reference associated with the first application, the one or more programs further including instructions for:

- determining whether the one or more words of the utterance are unrecognized by a fourth natural language model associated with a fourth application;
- marking the one or more unrecognized words as a second reference associated with the fourth application;
- determining whether a parameter of a second natural language model associated with a second application satisfies a requirement of the second reference;
- in accordance with a determination that the parameter of the second natural language model associated with the second application satisfies the requirement of the second reference, storing the relationship between the second reference and the parameter in the fourth natural language model associated with the fourth application.

14. The electronic device of claim 1, the one or more programs further including instructions for:

- in accordance with the determination that the parameter of the second natural language model associated with the second application satisfies the requirement of the reference, updating the first natural language model

67

associated with the first application to include the reference and the parameter of the second natural language model associated with the second application.

15. A method, comprising:

at an electronic device with one or more processors and memory:

receiving an utterance;

determining one or more words of the utterance unrecognized by a first natural language model associated with a first application;

marking the one or more unrecognized words as a reference;

determining whether a parameter of a second natural language model associated with a second application satisfies a requirement of the reference; and

in accordance with a determination that the parameter of the second natural language model associated with the second application satisfies the requirement of the reference, storing the relationship between the reference and the parameter in the first natural language model associated with the first application.

16. The method of claim **15**, wherein the utterance is received from a user.

17. The method of claim **15**, wherein the first natural language model was previously created by augmenting one or more utterances included in training data for the first application.

18. The method of claim **15**, wherein the second natural language model was previously created by augmenting one or more utterances included in training data for the second application.

19. The method of claim **15**, wherein determining whether the parameter of the second natural language model associated with the second application satisfies the requirement of the reference further comprises:

determining whether the parameter matches one or more properties of the reference.

20. The method of claim **19**, wherein the one or more properties of the reference includes at least one of a parameter type, a parameter value, and an application type.

21. The method of claim **15**, wherein storing the relationship between the reference and the parameter in the first natural language model associated with the first application further comprises:

updating the first natural language model associated with the first application to recognize the one or more unrecognized words as a reference to the parameter of the second natural language model.

22. The method of claim **17**, wherein augmenting one or more utterances included in training data for the first application further comprises changing one or more words of each of one of the utterances to a related word.

23. The method of claim **15**, wherein determining one or more words of the utterance unrecognized by a first natural language model associated with a first application further comprises:

mapping the one or more words of the utterance to an action structure associated with the first natural language model.

24. The method of claim **23**, wherein the action structure associated with the first natural language model comprises a template including one or more parameters.

25. The method of claim **15**, wherein the utterance is a first utterance and the reference is a first reference, further comprising:

68

receiving a second utterance;

determining whether the second utterance includes the first reference;

in accordance with a determination that the second utterance includes the first reference, replacing the first reference with the parameter of the second natural language model; and

in accordance with a determination that the second utterance does not include the reference:

determining one or more words of the second utterance unrecognized by a first natural language model associated with a first application; and

marking the one or more unrecognized words as a second reference.

26. The method of claim **15**, further comprising:

determining whether a parameter of a third natural language model associated with a third application satisfies a requirement of the reference;

in accordance with a determination that the parameter of the third natural language model associated with the third application satisfies the requirement of the reference, storing the relationship between the reference and the parameter in the first natural language model associated with the first application.

27. The method of claim **15**, wherein the reference is a first reference associated with the first application, further comprising:

determining whether the one or more words of the utterance are unrecognized by a fourth natural language model associated with a fourth application;

marking the one or more unrecognized words as a second reference associated with the fourth application;

determining whether a parameter of a second natural language model associated with a second application satisfies a requirement of the second reference;

in accordance with a determination that the parameter of the second natural language model associated with the second application satisfies the requirement of the second reference, storing the relationship between the second reference and the parameter in the fourth natural language model associated with the fourth application.

28. The method of claim **15**, further comprising:

in accordance with the determination that the parameter of the second natural language model associated with the second application satisfies the requirement of the reference, updating the first natural language model associated with the first application to include the reference and the parameter of the second natural language model associated with the second application.

29. A non-transitory computer-readable storage medium storing one or more programs, the one or more programs comprising instructions, which when executed by one or more processors of an electronic device, cause the electronic device to:

receive an utterance;

determine one or more words of the utterance unrecognized by a first natural language model associated with a first application;

mark the one or more unrecognized words as a reference;

determine whether a parameter of a second natural language model associated with a second application satisfies a requirement of the reference;

in accordance with a determination that the parameter of the second natural language model associated with the second application satisfies the requirement of the reference, store the relationship between the

reference and the parameter in the first natural language model associated with the first application.

30. The non-transitory computer-readable storage medium of claim 15, wherein the utterance is received from a user.

31. The non-transitory computer-readable storage medium of claim 15, wherein the first natural language model was previously created by augmenting one or more utterances included in training data for the first application.

32. The non-transitory computer-readable storage medium of claim 15, wherein the second natural language model was previously created by augmenting one or more utterances included in training data for the second application.

33. The non-transitory computer-readable storage medium of claim 15, wherein determining whether the parameter of the second natural language model associated with the second application satisfies the requirement of the reference further comprises:

determining whether the parameter matches one or more properties of the reference.

34. The non-transitory computer-readable storage medium of claim 33, wherein the one or more properties of the reference includes at least one of a parameter type, a parameter value, and an application type.

35. The non-transitory computer-readable storage medium of claim 29, wherein storing the relationship between the reference and the parameter in the first natural language model associated with the first application further comprises:

updating the first natural language model associated with the first application to recognize the one or more unrecognized words as a reference to the parameter of the second natural language model.

36. The non-transitory computer-readable storage medium of claim 31, wherein augmenting one or more utterances included in training data for the first application further comprises changing one or more words of each of one of the utterances to a related word.

37. The non-transitory computer-readable storage medium of claim 29, wherein determining one or more words of the utterance unrecognized by a first natural language model associated with a first application further comprises:

mapping the one or more words of the utterance to an action structure associated with the first natural language model.

38. The non-transitory computer-readable storage medium of claim 37, wherein the action structure associated with the first natural language model comprises a template including one or more parameters.

39. The non-transitory computer-readable storage medium of claim 15, wherein the utterance is a first utterance and the reference is a first reference, the one or more programs further comprising instructions, which when executed by one or more processors of an electronic device, cause the electronic device to:

receive a second utterance;
determine whether the second utterance includes the first reference;

in accordance with a determination that the second utterance includes the first reference, replace the first reference with the parameter of the second natural language model; and

in accordance with a determination that the second utterance does not include the reference:

determine one or more words of the second utterance unrecognized by a first natural language model associated with a first application; and

mark the one or more unrecognized words as a second reference.

40. The non-transitory computer-readable storage medium of claim 15, the one or more programs further comprising instructions, which when executed by one or more processors of an electronic device, cause the electronic device to:

determine whether a parameter of a third natural language model associated with a third application satisfies a requirement of the reference;

in accordance with a determination that the parameter of the third natural language model associated with the third application satisfies the requirement of the reference, store the relationship between the reference and the parameter in the first natural language model associated with the first application.

41. The non-transitory computer-readable storage medium of claim 15, wherein the reference is a first reference associated with the first application, the one or more programs further comprising instructions, which when executed by one or more processors of an electronic device, cause the electronic device to:

determine whether the one or more words of the utterance are unrecognized by a fourth natural language model associated with a fourth application;

mark the one or more unrecognized words as a second reference associated with the fourth application;

determine whether a parameter of a second natural language model associated with a second application satisfies a requirement of the second reference;

in accordance with a determination that the parameter of the second natural language model associated with the second application satisfies the requirement of the second reference, store the relationship between the second reference and the parameter in the fourth natural language model associated with the fourth application.

42. The non-transitory computer-readable storage medium of claim 15, the one or more programs further comprising instructions, which when executed by one or more processors of an electronic device, cause the electronic device to:

in accordance with the determination that the parameter of the second natural language model associated with the second application satisfies the requirement of the reference, update the first natural language model associated with the first application to include the reference and the parameter of the second natural language model associated with the second application.